Successive Concave Sparsity Approximation for Compressed Sensing

Mohammadreza Malek-Mohammadi*, *Member, IEEE*, Ali Koochakzadeh, Massoud Babaie-Zadeh, *Senior Member, IEEE*, Magnus Jansson, *Member, IEEE*, and Cristian R. Rojas, *Member, IEEE*

arXiv:1505.06841v2 [cs.IT] 26 Apr 2016

***Abstract*—In this paper, based on a successively accuracy-increasing approximation of the ℓ_0 norm, we propose a new algorithm for recovery of sparse vectors from underdetermined measurements. The approximations are realized with a certain class of concave functions that aggressively induce sparsity and their closeness to the ℓ_0 norm can be controlled. We prove that the series of the approximations asymptotically coincides with the ℓ_1 and ℓ_0 norms when the approximation accuracy changes from the worst fitting to the best fitting. When measurements are noise-free, an optimization scheme is proposed which leads to a number of weighted ℓ_1 minimization programs, whereas, in the presence of noise, we propose two iterative thresholding methods that are computationally appealing. A convergence guarantee for the iterative thresholding method is provided, and, for a particular function in the class of the approximating functions, we derive the closed-form thresholding operator. We further present some theoretical analyses via the restricted isometry, null space, and spherical section properties. Our extensive numerical simulations indicate that the proposed algorithm closely follows the performance of the oracle estimator for a range of sparsity levels wider than those of the state-of-the-art algorithms.**

***Index Terms*—Compressed Sensing (CS), Nonconvex Optimization, Iterative Thresholding, the LASSO Estimator, Oracle Estimator.**

EDICS Category: DSP-CPSL, DSP-RECO, SAM-CSSM, OPT-NCVX

I. Introduction

DURING the last decade with the appearance of theoretical and experimental results showing the possibility of recovering sparse signals from undersampled measurements, there has been an explosion of applications gaining from the reduction in the sampling rate. In fact, any field of science or engineering, where sampling a signal is part of the processing task, can potentially benefit from the line of research ongoing in the domain of compressed sensing (CS). This can be witnessed with applications in sciences such as quantum state tomography in physics [1], imaging in astrophysics [2], bacterial community reconstruction in biology [3], and genotyping in genetics [4]. In engineering, the number of applications is tremendous and includes magnetic resonance imaging [5], sampling of analog signals [6], array signal processing [7], and radar [8], to name a few.

A general formulation for the sampling in the compressed fashion is

$$\mathbf{b} = \mathbf{A}\widetilde{\mathbf{x}} + \mathbf{w}, \tag{1}$$

in which $\mathbf{b} \in \mathbb{R}^n$ is the vector of measurements, $\mathbf{A} \in \mathbb{R}^{n\times m}$ $(n < m)$ is the sensing matrix, $\widetilde{\mathbf{x}} \in \mathbb{R}^m$ is the unknown sparse vector with $s \ll m$ nonzero elements, and $\mathbf{w}$ is the probable additive noise. In a nutshell, given $\mathbf{b}$ and knowing $\mathbf{A}$, the goal in recovery of sparse vector from compressed measurements is to accurately estimate $\widetilde{\mathbf{x}}$. As (1) is underdetermined, recovery of $\widetilde{\mathbf{x}}$ from $\mathbf{b}$ is ill-posed unless we know *a priori* that $\widetilde{\mathbf{x}}$ resides in a low-dimensional space. Sparsity enters the game, and the sparsest solution which is consistent to the measurements is sought via

$$\min_{\mathbf{x}} \|\mathbf{x}\|_0 \quad \text{subject to} \quad \|\mathbf{A}\mathbf{x} - \mathbf{b}\| \leq \epsilon \tag{2}$$

in which $\|\mathbf{x}\|_0$ designates the so-called ℓ_0 norm of $\mathbf{x}$ defined as its number of nonzero entries, $\|\cdot\|$ stands for the ℓ_2 norm, and $\epsilon \geq \|\mathbf{w}\|$ is some constant.

Under certain circumstances, the solution to (2) is close to $\widetilde{\mathbf{x}}$ in (1) [9], [10], and, particularly, when $\mathbf{w} = \mathbf{0}$ (i.e., measurements are noise-free),

$$\min_{\mathbf{x}} \|\mathbf{x}\|_0 \quad \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b} \tag{3}$$

exactly recovers $\widetilde{\mathbf{x}}$ [9]. Programs (2) and (3) are generally NP-hard [11]; however, their convex relaxations,

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \quad \text{subject to} \quad \|\mathbf{A}\mathbf{x} - \mathbf{b}\| \leq \epsilon \tag{4}$$

and

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \quad \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b}, \tag{5}$$

can be considered instead. Convexification makes the recovery tractable at the cost of increasing the sampling rate needed for stable recovery along with increase in the reconstruction error [9]. A hot topic of research is, therefore, to propose recovery algorithms to push the sampling rate and reconstruction error as much as practically possible toward the intrinsic bounds of ℓ_0 minimization. To this end, several nonconvex alternatives for (3) and (2) have been proposed in the literature. The list for the noise-free recovery is already rich (cf. [12]–[16]), yet, in the noisy recovery, there is still much room for improvement.

This work was supported by the Swedish Research Council under contract 621-2011-5847 and the Swedish Strategic Research Area ICT-TNG program.

M. Malek-Mohammadi, M. Jansson and C. R. Rojas are with the ACCESS Linnaeus Centre, KTH (Royal Institute of Technology), Stockholm, 10044, Sweden (e-mail: {mohamma,janssonm,crro}@kth.se). A. Koochakzadeh is with the Department of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742 USA (email: alik@umd.edu). M. Babaie-Zadeh is with the Department of Electrical Engineering, Sharif University of Technology, Tehran 1458889694, Iran (e-mail: mbzadeh@yahoo.com).

A. Contribution

In this paper, we propose a new algorithm for recovery of sparse vectors. Although, as will be shown in the numerical simulations, the proposed algorithm outperforms some of the state-of-the-art algorithms in the noise-free recovery, the main goal of introducing this method is to provide an effective algorithm for the more challenging problem of recovery from compressed noisy measurements. The core idea of the proposed method is to closely approximate the ℓ_0 norm with a certain class of concave functions and then to minimize this approximation subject to the constraints. However, in contrast to the ideas of [13] and [12] where a fixed approximation is used, we use a series of approximations in which the accuracy of the approximations improves as the algorithm proceeds. The proposed algorithm is, hence, called SCSA standing for successive concave sparsity approximation.

In the noise-free case, our proposed algorithm involves solving some weighted ℓ_1 minimization programs which is computationally demanding. Nonetheless, in the noisy case, we utilize an iterative thresholding method [17] to decrease the complexity and derive a closed-form solution for the thresholding operator. In addition, we theoretically characterize the conditions under which the proposed iterative thresholding method converges.

On the theoretical side, the SCSA algorithm is supported by some guarantees derived from the restricted isometry [18], null space [19], and spherical section [20] properties. On the numerical side, extensive empirical evaluations in the noisy case show that the SCSA algorithm significantly and consistently outperforms ℓ_1 minimization as well as some other nonconvex methods in terms of reconstruction error, whereas the execution time is at most three times longer than that of one of the fastest algorithms in the comparison. Furthermore, we show that the SCSA algorithm closely follows the oracle estimator [21] for a broad range of sparsity levels.

B. Connections to Previous Work

The idea of the SCSA algorithm is borrowed from [22] that proposes a method for low-rank matrix recovery. In this work, we apply the same idea to the sparse recovery problem and propose an efficient optimization method for the noisy case which is not considered in [22]. Also, the underlying idea of SCSA is somehow connected to the ideas of the SL0 algorithm [14] and the algorithm of [23]. The fundamental difference with the SL0 algorithm is as follows. Let σ denote the parameter that reflects the accuracy of ℓ_0 norm approximation for SCSA, SL0, and the method of [23], where a smaller σ corresponds to better accuracy. For SCSA, we use a class of subadditive functions which enables us to analytically prove that, under some conditions, minimizing the ℓ_0 norm approximation for any $\sigma > 0$ leads to exact or accurate recovery in the noiseless or noisy cases. However, there is not such a guarantee for SL0 except for the asymptotic case of $\sigma \to 0$. This may also justify the performance improvement over SL0 demonstrated in our numerical simulations. Additionally, as the approximating functions differ, completely dissimilar optimization methods are exploited. In comparison to [23], the main distinctions are summarized as:

- The method of [23] has been proposed for the noise-free case and has been specialized for "reconstruction of sparse images," while SCSA applies to the noisy case as well and is designed for the general framework of CS.
- In this work, theoretical analyses for the asymptotic cases of $\sigma \to \infty$ and $\sigma \to 0$ are provided.
- Here, convergence analysis for the proposed optimization methods are also given.
- [23] solves the associated optimization problem by smoothing the approximating functions to make it simpler, while SCSA introduces no smoothing.
- Finally, in this paper, the proposed initialization is theoretically motivated, while [23] intuitively justifies its different initialization.

C. Notations and Outline

$\lceil x \rceil$ designates the smallest integer greater than or equal to x. $\mathrm{sign}(x) = |x|/x$ for $x \in \mathbb{R} \setminus \{0\}$, and $\mathrm{sign}(0) = 0$. The $\mathrm{sign}(\cdot)$, $\max(\cdot,\cdot)$ and $\min(\cdot,\cdot)$ functions act component-wise on vector inputs. For a vector $\mathbf{x}$, $\|\mathbf{x}\|_1$ and $\|\mathbf{x}\|$ denote the ℓ_1 and ℓ_2 norms, respectively, $\|\mathbf{x}\|_0$ denotes the so-called ℓ_0 norm, and x_i represents the ith element. For vectors $\mathbf{x}$ and $\mathbf{y}$, $\mathbf{x} \geq \mathbf{y}$ means that $\forall i, x_i \geq y_i$, $\mathbf{x} \odot \mathbf{y}$ indicates component-wise multiplication, and $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^T\mathbf{y}$ denotes the inner product. $\mathbf{X}^\dagger$ denotes the Moore-Penrose pseudoinverse of the matrix $\mathbf{X}$. For symmetric matrices $\mathbf{Y}, \mathbf{Z}$, $\mathbf{Y} \succeq \mathbf{Z}$ means $\mathbf{Y} - \mathbf{Z}$ is positive semidefinite. For a positive semidefinite matrix $\mathbf{X}$, $\lambda_{max}(\mathbf{X})$ and $\lambda_{min}(\mathbf{X})$ denote the largest and smallest eigenvalues. $\mathrm{null}(\mathbf{A})$ represents the null space of the matrix $\mathbf{A}$. $N(\mathbf{0}, \mathbf{\Sigma})$ designates the multivariate Gaussian distribution with mean $\mathbf{0}$ and covariance $\mathbf{\Sigma}$.

The rest of this paper is structured as follows. In Section II, the main idea of the SCSA algorithm is introduced, and Section III explains the optimization methods used in the noiseless and noisy cases. In Section IV, theoretical analyses are presented, and, in Section V, the performance of the SCSA method is evaluated and compared against some state-of-the-art algorithms by numerical simulations. Section VI concludes the paper.

II. Main idea of the SCSA Algorithm

A. Motivation

The problem of finding the sparsest solution of $\mathbf{Ax} = \mathbf{b}$ or the sparsest vector in the set $\{\mathbf{x} \mid \|\mathbf{Ax} - \mathbf{b}\| \leq \epsilon\}$ can be interpreted as the task of approximating the Kronecker delta function (called herein the delta function for brevity). Let

$$\delta(x) = \begin{cases} 1 & \text{if } x = 0 \\ 0 & \text{otherwise} \end{cases}$$

denote the delta function, then the ℓ_0 norm of a vector $\mathbf{x}$ is equal to $\|\mathbf{x}\|_0 = \sum_{i=1}^{m}[1 - \delta(x_i)]$ and can be approximated by $\sum_{i=1}^{m} f(x_i)$ in which $f : \mathbb{R} \to \mathbb{R}$ acts as a delta approximating (DA) function.[1] Replacing the ℓ_0 norm with

[1]Clearly, $f(x)$ is approximating $1 - \delta(x)$, not the delta function. Nevertheless, we call it DA function for the sake of easy referral.

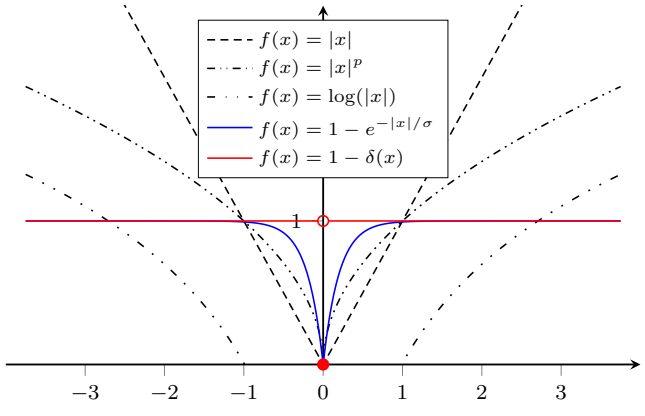


Fig. 1. Some algorithms for recovery of sparse vectors implicitly or explicitly use different functions to approximate the delta function as $\|\mathbf{x}\|_0$ equals $\sum_{i=1}^{n} 1 - \delta(x_i)$. Some of them are plotted in this figure, and among them, $f_\sigma(x) = 1 - e^{-|x|/\sigma}$ more closely matches $1 - \delta(x)$ for a small σ.

the above approximation, the next step is to find a point in the feasible set that minimizes $\sum_{i=1}^{m} f(x_i)$. To have a numerically tractable optimization problem, one needs to find suitable DA functions with some appropriate properties like convexity or continuity. Based on this formulation, $f(x) = |x|$ promotes sparsity with the ℓ_1 norm, and $f(x) = |x|^p, 0 < p < 1$, gives rise to ℓ_p quasi-norm minimization. Another well-known example is $\log(x)$ which leads to reweighted ℓ_1 minimization [12].

Though ℓ_1 minimization enjoys convexity, a number of theoretical and numerical analyses show the superiority of other DA choices. For instance, [24] and [25] show that, at least for some p's in $(0, 1)$ and under smaller thresholds on the restricted isometry constant (RIC) [18], globally minimizing the ℓ_p quasi-norm subject to the linear equations recovers sparse vectors uniquely. Also, [26], [27] prove that, in noisy CS and under some mild conditions, even locally minimizing the ℓ_p quasi-norm via a certain optimization scheme leads to recovery of a solution with less error. On the numerical side, [13], [24], [28] give some numerical evaluations demonstrating the superiority of ℓ_p quasi-norm minimization. Moreover, [12], [29], [30] analyze and compare the theoretical and/or numerical performance of reweighted ℓ_1 minimization to ℓ_1 minimization showing considerable improvements.

With this background, one may think whether better approximations of the delta function lead to higher performance in recovery of sparse vectors. In this paper, we use a class of DA functions which more closely approximate the delta function and show that this intuition is indeed the case. Fig. 1 shows the above DA functions as well as $1 - \exp(-|x|/\sigma)$, one of the exploited DA functions in this paper in which σ is a parameter to control the fitting to the delta function. Obviously, by choosing a small enough σ, it is possible to have the best fit to the delta function among the plotted DA functions. Putting this DA function in the sparse recovery problem, we are proposing to solve

$$\min_{\mathbf{x}} \sum_{i=1}^{m} \Big(1 - \exp(-|x_i|/\sigma)\Big) \quad \text{s.t.} \quad \|\mathbf{Ax} - \mathbf{b}\| \le \epsilon, \qquad (6)$$

where ϵ is equal to 0 in the noise-free case, to obtain a solution to (3) or (2).

B. Properties of DA Functions

To establish theoretical analysis for the proposed algorithm and derive efficient optimization methods, we need to impose some assumptions on the DA functions, summarized in the following property.

Property 1: Let $f : \mathbb{R} \to [-\infty, \infty)$ and define $f_\sigma(x) \triangleq f(x/\sigma)$ for any $\sigma > 0$. The function f is said to possess Property 1, if

(a) f is real analytic on (x_0, ∞) for some $x_0 < 0$,
(b) $\forall x \ge 0$, $f''(x) \ge -c$, where $c > 0$ is some constant,
(c) f is concave on $\mathbb{R}$,
(d) $f(x) = 0 \Leftrightarrow x = 0$,
(e) $\lim_{x\to+\infty} f(x) = 1$.

It follows immediately from Property 1 that $\{f_\sigma(|x|)\}$ converges pointwise to $1 - \delta(x)$ as $\sigma \to 0^+$; i.e.,

$$\lim_{\sigma\to 0^+} f_\sigma(|x|) = \begin{cases} 0 & \text{if } x = 0 \\ 1 & \text{otherwise.} \end{cases}$$

Besides the DA function $f(x) = 1 - e^{-x}$ which is mainly used in this paper, there are other functions that satisfy conditions of Property 1 including

$$f(x) = \begin{cases} \dfrac{x}{x+1} & \text{if } x \ge x_0 \\ -\infty & \text{otherwise} \end{cases}$$

for some $-1 < x_0 < 0$.

C. The Main Idea

To obtain higher performance in recovery of sparse vectors, we propose to closely approximate the ℓ_0 norm and then solve the consequent optimization problem. More precisely, let $f(\cdot)$ denote a function possessing Property 1, then we define $F_\sigma(|\mathbf{x}|) = \sum_{i=1}^{m} f_\sigma(|x_i|) \approx \|\mathbf{x}\|_0$ and use

$$\min_{\mathbf{x}} \Big(F_\sigma(|\mathbf{x}|) = \sum_{i=1}^{m} f_\sigma(|x_i|)\Big) \quad \text{subject to} \quad \mathbf{Ax} = \mathbf{b} \qquad (7)$$

to find a solution to (3) and

$$\min_{\mathbf{x}} F_\sigma(|\mathbf{x}|) \quad \text{subject to} \quad \|\mathbf{Ax} - \mathbf{b}\| \le \epsilon \qquad (8)$$

to obtain a solution to (2). Expectedly, these optimization problems are not convex, and any algorithm may get stuck in a local minimum.

Intuitively, when σ is small and $F_\sigma(|\mathbf{x}|)$ approximates $\|\mathbf{x}\|_0$ with a good accuracy, there are many local solutions which makes the task of optimizing (7) and (8) very hard. In contrast, if σ is relatively large, while the accuracy of the approximation is not good, $F_\sigma(|\mathbf{x}|)$ has a smaller number of local minima. In line with this, in the asymptotic case, it will be shown that $\sigma F_\sigma(|\mathbf{x}|)$ becomes convex as σ tends to infinity.

Following the same approach as in [14], [23], [31], a continuation scheme for solving (7) and (8) is utilized which helps in achieving the sparsest solution of $\mathbf{Ax} = \mathbf{b}$ and

$\{\mathbf{x} \mid \|\mathbf{A}\mathbf{x}-\mathbf{b}\| \leq \epsilon\}$ instead of finding a local minimum. Initially, optimization of (7) or (8) is started with a very large value of σ, and the solution is passed as an initial guess to the next iteration in which (7) or (8) is solved for a smaller value of σ. These iterations continue until reaching a desired accuracy.

To further decrease the chance of getting trapped in local minima, we constrain $f_\sigma(\cdot)$ to be continuous with respect to σ in Property 1. In this fashion, we expect that when σ's at two consecutive iterations are close, the global minimizers for these two σ's are in the vicinity of each other. Thus, starting from a convex optimization and gradually decreasing σ, it is more likely that a global solution will be found.

III. Implementation of the SCSA algorithm

Contrary to [23] that tries to solve (7) by converting it to an unconstrained problem and smoothing the DA function around the origin, we solve (7) directly by employing a majorize-minimize (MM) technique [32] without smoothing the DA function. This approach leads to iteratively solving a few weighted ℓ_1 minimization (or linear) programs. Inspired by the iterative thresholding (IT) technique, we also propose two efficient methods for the noisy case of (8) which are computationally attractive.

A. Optimization for a Fixed σ in the Noise-Free Case

Although $F_\sigma(|\mathbf{x}|)$ is not a differentiable function, by restricting $\mathbf{x}$ to be in the positive orthant, we can drop $|\cdot|$ from its argument, and make it differentiable. In other words, if we look for the sparsest solution which belongs to the positive orthant, then we can use $f_\sigma(x) = 1-\exp(-x/\sigma)$ with the constraint $\mathbf{x} \geq 0$ in (6) and, in this way, make the cost function differentiable. Nevertheless, exploiting the same technique as in the conversion of (5) to a linear program [11], we can overcome this restriction. To be precise, let $\mathbf{y} = \begin{bmatrix}\mathbf{x}_p^T & \mathbf{x}_m^T\end{bmatrix}^T$ denote a column vector of length $2m$ in which $\mathbf{x}_p = \max(\mathbf{x}, \mathbf{0})$ and $\mathbf{x}_m = -\min(\mathbf{x}, \mathbf{0})$. The elements of $\mathbf{y}$ are nonnegative, $F_\sigma(\mathbf{y}) = F_\sigma(|\mathbf{x}|)$, and the constraints $\mathbf{A}\mathbf{x} = \mathbf{b}$ are converted to $[\mathbf{A}, -\mathbf{A}]\mathbf{y} = \mathbf{b}$. By introducing this new vector, (7) can be reformulated as

$$\min_{\mathbf{y}} F_\sigma(\mathbf{y}) \quad \text{subject to} \quad [\mathbf{A}, -\mathbf{A}]\mathbf{y} = \mathbf{b},\ \mathbf{y} \geq \mathbf{0}. \tag{9}$$

The following theorem proves that, by solving (9), we are able to optimize program (7).

Theorem 1: For any class of functions possessing Property 1, programs (7) and (9) are equivalent.

Proof: The proof follows the same lines as in the proof of equivalence of ℓ_1-minimization to a linear program [11]. Let $\mathbf{y}^* = \begin{bmatrix}\mathbf{u}^T\mathbf{v}^T\end{bmatrix}^T$, where $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$, denote an optimal solution to (9). To show that (7) and (9) are equivalent, it is sufficient to show that the definition of $\mathbf{y}$ in (9) is not violated, or, mathematically speaking, the supports of $\mathbf{u}$ and $\mathbf{v}$ do not overlap. Assume, to the contrary, that they overlap at index k. Without loss of generality, further assume that $u_k > v_k$, then there is another solution $\widehat{\mathbf{y}} = \begin{bmatrix}\widehat{\mathbf{u}}^T\widehat{\mathbf{v}}^T\end{bmatrix}^T$ with $\widehat{\mathbf{u}} = \mathbf{u}$ and $\widehat{\mathbf{v}} = \mathbf{v}$ except for $\widehat{u}_k = u_k - v_k$ and $\widehat{v}_k = 0$. This new solution is feasible since $\widehat{\mathbf{u}} \geq \mathbf{0}$, $\widehat{\mathbf{v}} \geq \mathbf{0}$, and $[\mathbf{A}, -\mathbf{A}]\widehat{\mathbf{y}} = \mathbf{b}$. However, $F_\sigma(\widehat{\mathbf{y}})$ is smaller than $F_\sigma(\mathbf{y}^*)$ by $f_\sigma(u_k) + f_\sigma(v_k) - f_\sigma(u_k - v_k) > 0$ which contradicts the optimality of $\mathbf{y}^*$. ■

Since (9) is a concave program, the MM technique can be easily used to find at least a local solution. First-order concavity condition for $F_\sigma(\cdot)$ implies that

$$F_\sigma(\mathbf{y}) \leq F_\sigma(\widetilde{\mathbf{y}}) + \langle \mathbf{y} - \widetilde{\mathbf{y}}, \nabla F_\sigma(\widetilde{\mathbf{y}})\rangle = H_\sigma(\mathbf{y}, \widetilde{\mathbf{y}})$$

for some $\widetilde{\mathbf{y}}$ in the feasible set. To apply the MM technique, $H_\sigma(\cdot,\cdot)$ is selected as a surrogate function. Consequently, neglecting the fixed terms, to obtain a solution to (7), one needs to iteratively solve

$$\mathbf{y}_{k+1} = \operatorname*{argmin}_{\mathbf{y}} \Big\{ \langle \nabla F_\sigma(\mathbf{y}_k), \mathbf{y}\rangle \ \Big|\ [\mathbf{A}, -\mathbf{A}]\mathbf{y} = \mathbf{b},\ \mathbf{y} \geq \mathbf{0} \Big\} \tag{10}$$

for $k \geq 0$ until convergence. It can be verified that the program (10) is equal to a weighted ℓ_1 minimization

$$\mathbf{x}_{k+1} = \operatorname*{argmin}_{\mathbf{x}} \Big\{ \|\mathbf{W}\mathbf{x}\|_1 \ \Big|\ \mathbf{A}\mathbf{x} = \mathbf{b} \Big\}, \tag{11}$$

where $\mathbf{W} = \operatorname{diag}(\nabla F_\sigma(\mathbf{z}))$ and $\mathbf{z} = |\mathbf{x}_k|$.

The next proposition, which can be easily deduced from [22, Theorem 3], proves the convergence of the proposed MM based approach.

Proposition 1: The sequence $\{\mathbf{y}_k\}$ obtained from (10) is convergent to a local minimum of (9), and $F_\sigma(\mathbf{y}_{k+1}) \leq F_\sigma(\mathbf{y}_k)$ for all $k \geq 0$.

B. Optimization for a Fixed σ in the Noisy Case

A computationally attractive way to find a solution to an unconstrained optimization problem, in which the cost function is composed of the sum of a smooth and a nonsmooth convex function, is to use iterative thresholding or proximal algorithms; see [17] for a comprehensive discussion. This kind of problems is represented as

$$\min_{\mathbf{x}} \lambda\rho(\mathbf{x}) + h(\mathbf{x}), \tag{12}$$

where $\rho(\mathbf{x})$ is convex but possibly nonsmooth, $h(\mathbf{x})$ is convex and differentiable with Lipschitz continuous gradient, and λ is a regularization parameter. Particularly, program (4) can be converted to an equivalent unconstrained optimization problem, known also as the LASSO program [33],

$$\min_{\mathbf{x}} \lambda\|\mathbf{x}\|_1 + \|\mathbf{A}\mathbf{x}-\mathbf{b}\|^2, \tag{13}$$

where $\lambda > 0$ is a constant to regularize between solution sparsity and consistency to measurements. Accordingly, (4) can be solved by applying iterative thresholding method on (13). This special case of proximal methods is known as the iterative soft thresholding (IST) method. More generally, (13) can be written as

$$\min_{\mathbf{x}} \lambda\|\mathbf{x}\|_1 + h(\mathbf{x}), \tag{14}$$

where $h(\mathbf{x})$ is as in (12). The IST method is aimed for finding a solution to (14) by iteratively solving [34]

$$\mathbf{x}_{k+1} = \operatorname*{argmin}_{\mathbf{x}} \Big\{ \langle \mathbf{x}-\mathbf{x}_k, \nabla h(\mathbf{x}_k)\rangle + \frac{1}{2\mu}\|\mathbf{x}-\mathbf{x}_k\|^2 + \lambda\|\mathbf{x}\|_1 \Big\},$$

where μ is a step-size parameter. The above program admits a unique closed-form solution given by

$$\mathbf{x}_{k+1} = \mathcal{S}_{\lambda\mu}(\mathbf{x}_k - \mu\nabla h(\mathbf{x}_k)), \tag{15}$$

where $\mathcal{S}_\alpha : \mathbb{R} \to \mathbb{R}$ is the soft thresholding operator defined as $\mathcal{S}_\alpha(x) = \max(|x| - \alpha, 0)\,\mathrm{sign}(x)$ for a scalar input and is applied component-wise to vectors.

To utilize the IT method for finding a solution to (8), first, program (8) is formulated as an unconstrained optimization problem

$$\min_{\mathbf{x}} \lambda_\sigma F_\sigma(|\mathbf{x}|) + h(\mathbf{x}), \tag{16}$$

where $h(\mathbf{x})$ is as in (12) and λ_σ is a regularization parameter similar to (14) and, in general, may be a function of σ.

Similar to the convex case, (16) can be optimized by iteratively solving

$$\mathbf{x}_{k+1} = \operatorname*{argmin}_{\mathbf{x}} \Big\{\langle \mathbf{x} - \mathbf{x}_k, \nabla h(\mathbf{x}_k)\rangle + \frac{1}{2\mu}\|\mathbf{x} - \mathbf{x}_k\|^2 + \lambda_\sigma F_\sigma(|\mathbf{x}|)\Big\} \tag{17}$$

until convergence. Ignoring constant terms, it can be verified that the program (17) is equal to

$$\mathbf{x}_{k+1} = \operatorname*{argmin}_{\mathbf{x}} \Big\{ \frac{1}{2\mu}\|\mathbf{x} - (\mathbf{x}_k - \mu\nabla h(\mathbf{x}_k))\|^2 + \lambda_\sigma F_\sigma(|\mathbf{x}|)\Big\}. \tag{18}$$

For a general $F_\sigma(\cdot)$, one should use iterative methods to solve (18). However, remarkably, for $f_\sigma(|x|) = 1 - \exp(-|x|/\sigma)$, we can find a closed-from solution using *the Lambert W* function [35] which enables us to efficiently solve (18). The details of derivation of the closed-form solution as well as the corresponding thresholding operator are given in Appendix A. Using this thresholding operator, for any fixed σ, our IT based approach for solving (16) simplifies to iteratively updating $\mathbf{x}_{k+1}$ by

$$\mathbf{x}_{k+1} = \mathcal{T}^{(\sigma)}_{\lambda_\sigma\mu}\big(\mathbf{x}_k - \mu\nabla h(\mathbf{x}_k)\big), \tag{19}$$

where $\mathcal{T}^{(\sigma)}_\alpha$ is given in (25) in Appendix A.

The next theorem whose proof is left for Appendix B analyzes the convergence of the sequence generated by (17) or (19).

Theorem 2: Let $M > 0$ denote the smallest Lipschitz constant of ∇h, and let M' be the smallest value such that $\lambda_\sigma\nabla^2 F_\sigma(\mathbf{x}) \succeq -M'\mathbf{I}, \forall \mathbf{x} \geq \mathbf{0}$.[2] For any $\mu \in (0, \frac{1}{M+M'})$, the sequence $\{\mathbf{x}_k\}$ generated by (17) is convergent to a stationary point of (16).

Although the IT technique has a low complexity, it lacks fast convergence rate [34]. To speed up the rate of convergence, a Nesterov's like acceleration step is added in [34] to the IST method which considerably increases the rate of convergence. While the complexity does not boost, improvement in the convergence rate is considerable. To accelerate our IT based approach, without proving the convergence, we also exploit a similar technique based on the FISTA algorithm [34]. We call this instance of the SCSA algorithm fast iterative thresholding (FIT) based.

[2]For $h(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$ and $F_\sigma(\mathbf{x}) = \sum_{i=1}^m 1 - e^{-x_i/\sigma}$, $M = 2\lambda_{max}(\mathbf{A}^T\mathbf{A})$ and $M' = \lambda_\sigma/\sigma^2$.

Now, let us remark on how the regularization parameter λ_σ should be scaled with σ. For (16) with $h(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$, a necessary condition for optimality of a point $\mathbf{x}^*$ is that

$$\mathbf{0} \in 2\mathbf{A}^T\mathbf{A}\mathbf{x}^* - 2\mathbf{A}^T\mathbf{b} + \lambda_\sigma\partial F_\sigma(|\mathbf{x}^*|) \tag{20}$$

where $\partial F_\sigma(|\mathbf{x}^*|)$ denotes the Clarke subdifferential [36] of $F_\sigma(|\cdot|)$ at point $\mathbf{x}^*$. Particularly, in the scalar case, $\partial f_\sigma(|x|)$ for $f_\sigma(x) = 1 - e^{-x/\sigma}$ is given by

$$\partial f_\sigma(|x|) = \begin{cases} \Big\{\dfrac{\mathrm{sign}(x)}{\sigma} e^{-\frac{|x|}{\sigma}}\Big\} & \text{if } x \neq 0 \\ [-1/\sigma, 1/\sigma] & \text{if } x = 0. \end{cases}$$

Let $\tau, \mathbf{x}^*_\tau$, and $\mathbf{A}_\tau$ denote the support set of $\mathbf{x}^*$ and restriction of $\mathbf{x}^*$ and $\mathbf{A}$ to the entries and columns indexed by τ, respectively. Condition (20) implies that

$$\mathbf{x}^*_\tau = \mathbf{A}^\dagger_\tau\mathbf{b} - \frac{\lambda_\sigma}{2\sigma}(\mathbf{A}^T_\tau\mathbf{A}_\tau)^{-1}\,\mathrm{sign}(\mathbf{x}^*_\tau) \odot e^{-\frac{|\mathbf{x}^*_\tau|}{\sigma}}, \tag{21}$$

$$|\mathbf{a}_i^T(\mathbf{A}_\tau\mathbf{x}^*_\tau - \mathbf{b})| \leq \frac{\lambda_\sigma}{2\sigma}, \quad i \notin \tau, \tag{22}$$

where $\mathbf{a}_i$ denotes the ith column of $\mathbf{A}$. The second term in the right-hand side of (21) disappears when $\sigma \to 0$, since the exponential terms decay faster than σ. Therefore, if τ coincides with the support set of the true solution, (21) shows that $\mathbf{x}^*_\tau$ tends to the oracle solution, which is obtained by knowing the support set of the true solution. However, when σ is decreased to smaller values, inequality (22) will be satisfied easier, and (16) will give solutions with smaller sparsity levels. In fact, $|\mathbf{a}_i^T(\mathbf{A}_\tau\mathbf{x}^*_\tau - \mathbf{b})|$ measures how close are the residual $\mathbf{A}_\tau\mathbf{x}^*_\tau - \mathbf{b}$ and the ith column of A; hence, the larger the threshold, the larger the number of indices of columns of $\mathbf{A}$ to be excluded from τ. Moreover, the threshold λ_σ/σ should naturally be independent of σ but dependent on the noise level. Thus, to have this threshold independent of σ, λ_σ is scaled linearly with σ; i.e., $\lambda_\sigma = \lambda\sigma$ for some constant λ.

A closer look at the thresholding operator introduced in (25) reveals an interesting interpolation property. Formally, in Propositions 2 and 4, we will show how (7) asymptotically converges to ℓ_0 and ℓ_1 minimization when σ goes to 0 or ∞. The thresholding operator $\mathcal{T}^{(\sigma)}_{\lambda_\sigma}(\cdot)$, however, also more simply illustrates this asymptotical behavior for $1 - \exp(-|x|/\sigma)$. Figure 2 displays $\mathcal{T}^{(\sigma)}_{\lambda_\sigma}(\cdot)$ for $\sigma = 100, 1$, and 0.1 when $\lambda_\sigma = \lambda\sigma$ and λ is fixed to 1. In this plot, when σ is relatively large, $\mathcal{T}^{(\sigma)}_{\lambda_\sigma}(\cdot)$ is very close to the soft thresholding operator, whereas, for a small σ, it is very similar to the hard thresholding operator [37], which according to the formulation used in this paper is defined as

$$\mathcal{H}_\lambda(x) = |x|\mathbb{1}_{|x|>\sqrt{2\lambda}}, \tag{23}$$

where $\mathbb{1}$ denotes the indicator function. This shows that when σ is swept from very large to smaller values, the thresholding operator gradually converts from the soft thresholding operator to the hard thresholding operator, making an interpolation between ℓ_1 and ℓ_0 minimization.

Finally, it is worth comparing our above proposed approach to a few other available methods. In [26] and [38], a multi-stage convex relaxation method for solving (16) has been

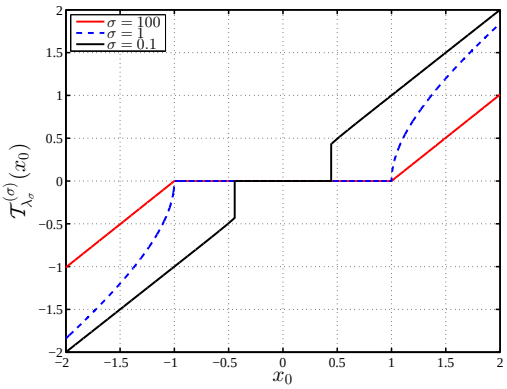


Fig. 2. $\mathcal{T}^{(\sigma)}_{\lambda_\sigma}$ is plotted for $\sigma = 100, 1,$ and 0.1 with $\lambda = 1$ and $\lambda_\sigma = \lambda\sigma$. For the large σ, $\mathcal{T}^{(\sigma)}_{\lambda_\sigma}$ is very close to the soft thresholing operator $\mathcal{S}_\lambda$ with $\lambda = 1$, whereas it is very similar to the hard thresholding operator defined in (23) with $\lambda = 0.1$.

proposed, where the nonconvex function $F(\cdot)$ is more general than those defined by Property 1.[3] When $F(\cdot)$ is aimed for sparsity regularization, the proposed optimization method involves solving a number of weighted versions of (13). Namely, one needs to iteratively solve

$$\mathbf{x}_{k+1} = \underset{\mathbf{x}}{\operatorname{argmin}} \Big\{ \lambda \|\mathbf{W}_k \mathbf{x}\|_1 + \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \Big\},$$

where $h(\mathbf{x})$ is assumed to be equal to $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$ and $\mathbf{W}_k$ is the weighting matrix that depends on $\mathbf{x}_k$ (the previous solution) and the function $F(\cdot)$ [26], [38]. In contrast to this approach and instead of solving a sequence of optimization problems to obtain a solution to (16), our proposed method directly solves (16) which in turn seems to be more computationally efficient. Moreover, in [39], a regularization path-following algorithm for solving (16) is introduced which is based on the iterative thresholding approach. Along with this algorithm, some strong theoretical guarantees are provided for a special class of nonconvex regularizations. This algorithm as well as its theoretical analyses, however, cannot be applied to the DA functions considered herein since, for instance, $f_\sigma(|x|) = 1 - \exp(-|x|/\sigma)$ does not satisfy regularity condition (a) in [39].

C. Initialization

The following proposition proves that when $\sigma \to \infty$, a scaled version of $F_\sigma(|\mathbf{x}|)$ becomes equal to the ℓ_1 norm. Consequently, the MM based instance of the proposed algorithm is initialized with a minimum ℓ_1-norm solution of $\mathbf{A}\mathbf{x} = \mathbf{b}$, and the IT and FIT based instances are initialized with the solution obtained by the FISTA algorithm. The proof easily follows from [22, Theorem 1].

Proposition 2: For any class of functions $\{f_\sigma\}$ possessing Property 1,

$$\lim_{\sigma \to \infty} \frac{\sigma}{\gamma} F_\sigma(|\mathbf{x}|) = \|\mathbf{x}\|_1,$$

where $\gamma = f'(0) \neq 0$.

[3] In the formulation of [26], $F(\cdot)$ also does not depend on any scaling parameter like σ.

Subalgorithm 1 Opt_FIT: Optimization of (16) for a fixed σ using the FISTA like acceleration.

Input: $\mathbf{A}, \mathbf{b}, \lambda, \mu, \sigma, \mathbf{x}_0,$ `tol`
Initialization:
$\mathbf{y}_1 = \mathbf{x}_0,\ t_1 = 1,\ k = 0,\ h(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2.$
Body:
while $d >$ `tol` **do**
 $k = k + 1.$
 $\mathbf{x}_k = \mathcal{T}^{(\sigma)}_{\lambda\mu}(\mathbf{y}_k - \mu\nabla h(\mathbf{y}_k)).$
 $t_{k+1} = (1 + \sqrt{1 + 4t_k})/2.$
 $\mathbf{y}_{k+1} = \mathbf{x}_k + (t_k - 1)(\mathbf{x}_k - \mathbf{x}_{k-1})/t_{k+1}.$
 $d = \|\mathbf{x}_k - \mathbf{x}_{k-1}\|/\|\mathbf{x}_{k-1}\|.$
end while
Output: $\mathbf{x}_k$

Algorithm 1 The SCSA algorithm

Input: $\mathbf{A}, \mathbf{b}, \lambda$ (Only for IT and FIT based instances)
Initialization:
$c \in (0, 0.5)$: decreasing factor for σ.
$\epsilon_1, \epsilon_2 > 0$: stopping thresholds for main and internal loops.
$h(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2, f_\sigma(x) = 1 - e^{-|x|/\sigma}.$
$$\begin{cases} \text{LP based:} & \mathbf{x}_0 = \operatorname{argmin}_{\mathbf{x}} \big\{ \|\mathbf{x}\|_1 | \mathbf{A}\mathbf{x} = \mathbf{b} \big\}. \\ \text{IT or FIT based:} & \mathbf{x}_0 = \operatorname{argmin}_{\mathbf{x}} \big\{ \lambda\|\mathbf{x}\|_1 + \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 \big\}. \end{cases}$$
$\sigma_0 = 8\max(|\mathbf{x}_0|).$
Body:
$i = 0, \sigma = \sigma_0.$
while $d_1 > \epsilon_1$ **do**
 $\widehat{\mathbf{x}}_0 = \mathbf{x}_i, j = 0, i = i + 1, \lambda_\sigma = \lambda\sigma.$
 while $d_2 > \epsilon_2$ **do**
 $j = j + 1.$
 $\mathbf{W} = \operatorname{diag}(\nabla F_\sigma(|\widehat{\mathbf{x}}_{j-1}|)).$
 $$\begin{cases} \text{LP based:} & \widehat{\mathbf{x}}_j = \operatorname{argmin}_{\mathbf{x}} \big\{ \|\mathbf{W}\mathbf{x}\|_1 \ \big|\ \mathbf{A}\mathbf{x} = \mathbf{b} \big\}. \\ \text{IT based:} & \widehat{\mathbf{x}}_j = \mathcal{T}^{(\sigma)}_{\lambda_\sigma \mu}\big(\widehat{\mathbf{x}}_{j-1} - \mu\nabla h(\widehat{\mathbf{x}}_{j-1})\big). \\ \text{FIT based:} & \widehat{\mathbf{x}}_j = \text{Opt_FIT}(\mathbf{A}, \mathbf{b}, \lambda_\sigma, \mu, \sigma, \widehat{\mathbf{x}}_{j-1}, \epsilon_2). \end{cases}$$
 $d_2 = \|\widehat{\mathbf{x}}_j - \widehat{\mathbf{x}}_{j-1}\|/\|\widehat{\mathbf{x}}_{j-1}\|.$
 end while
 $\mathbf{x}_i = \widehat{\mathbf{x}}_j.$
 $d_1 = \|\mathbf{x}_i - \mathbf{x}_{i-1}\|/\|\mathbf{x}_{i-1}\|.$
 $\sigma = c\sigma.$
end while
Output: $\mathbf{x}_i$

D. The Final Algorithm

Putting all the above steps together, the final algorithm, summarized in Algorithm 1, is obtained by exploiting $f_\sigma(x) = 1 - e^{-x/\sigma}$ and $h(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$. According to the method used in optimizing (7) or (16) for a fixed σ, three instances of the SCSA algorithm are summarized as

- SCSA-LP (based on linear programming),
- SCSA-IT (based on IT method),
- SCSA-FIT (based on FIT method).

To completely characterize the implementation of this algorithm, the following remarks are in order.

Remark 1. Parameter σ is decayed by a multiplicative factor c which should be chosen in the interval $(0, 0.5)$. We will discuss how to properly choose it in Section V with more details. Furthermore, following the same reasoning as in [22], σ_0 is set to $8\max(|\mathbf{x}_0|)$ because this σ_0 virtually acts as if σ

tends to ∞.

Remark 2. $d_1 = \|\mathbf{x}_i - \mathbf{x}_{i-1}\|/\|\mathbf{x}_{i-1}\|$ and $d_2 = \|\widehat{\mathbf{x}}_j - \widehat{\mathbf{x}}_{j-1}\|/\|\widehat{\mathbf{x}}_{j-1}\|$ measure relative distances between the solution of successive iterations of the external and internal loops, respectively, and are used to stop execution of these loops. In the proposed continuation approach, it is not necessary to run the internal loop until convergence, and it is just needed to get close to the minimizer for the current value of σ. Consequently, ϵ_1 is usually set a few orders of magnitude smaller than ϵ_2. For the noise-free case, $\epsilon_1 = 10^{-3}$ and $\epsilon_2 = 10^{-2}$ are suggested. In the noisy case, since the IT and FIT based approaches are relatively slow and the difference between two consecutive solutions is not large, ϵ_1 and ϵ_2 should be chosen smaller. Moreover, as SCSA-IT generally has a slower convergence rate, ϵ_2 (the threshold for the internal loop) for this instance of SCSA should be smaller than that of SCSA-FIT. This choice, as will be shown in numerical simulations, leads to a similar performance in terms of reconstruction accuracy. ϵ_1 and ϵ_2 are also functions of the regularization parameter λ. Putting altogether, we numerically found that a good choice for SCSA-IT is $\epsilon_1 = \min(10^{-4}, 10^{-3}\lambda)$ and $\epsilon_2 = \min(10^{-4}, 10^{-3}\lambda)$ and for SCSA-FIT is $\epsilon_1 = \min(10^{-4}, 10^{-3}\lambda)$ and $\epsilon_2 = \min(10^{-3}, 10^{-2}\lambda)$.

Remark 3. Proposition 2 can be strengthened to

$$\lim_{\sigma\to\infty} \operatorname*{argmin}_{\mathbf{x}} \left\{F_\sigma(|\mathbf{x}|)|\mathbf{A}\mathbf{x} = \mathbf{b}\right\} = \operatorname*{argmin}_{\mathbf{x}} \left\{\|\mathbf{x}\|_1|\mathbf{A}\mathbf{x} = \mathbf{b}\right\}$$

provided that the above ℓ_1 minimization has a unique solution. Nevertheless, since not a strictly convex program, it may occur that ℓ_1 minimization does not admit a unique solution. Let S_0 and S_1 denote the solution sets of (3) and (5), respectively. An interesting problem is to characterize the conditions under which S_0 (which we assume to be singleton) is a subset of S_1. Given these conditions, one can hope to devise a suitable optimization algorithm which is theoretically guaranteed to start from a point in S_1 and end up in the unique solution of (3). In this fashion, the guaranteed recovery bounds for ℓ_1 minimization can be improved.

Remark 4. As discussed earlier, λ_σ is set to $\lambda\sigma$ for some $\lambda > 0$. This choice can be also justified as follows. Since SCSA in the noisy case is initialized with a solution to (13) and $\lim_{\sigma\to\infty} \sigma F_\sigma(|\mathbf{x}|) = \|\mathbf{x}\|_1$, it is natural to set $\lambda_\sigma = \lambda\sigma$ to have the same cost function in (13) and (16) for $\sigma \to \infty$.

IV. Theoretical Analysis

A thorough performance analysis of the SCSA algorithm considering all of its steps seems to be very hard and cannot be embedded in this paper. In this section, however, by analyzing programs (7) and (8) for any $\sigma > 0$ and/or σ tending to 0^+, we provide simplified analyses which will give the reader a theoretical insight about how the main idea works. These analyses are simply extracted from the results in [19], [22], [40], [41] and are based on the null-space [19], restricted isometry [18], and spherical section properties [20], [42] of the sensing matrix. We recall or modify them in order to be able to study the performance of the proposed algorithm.

Null-space based recovery conditions: It can be verified that Property 1 implies the 'sparseness measure' definition in [19]. Consequently, based on Theorems 2 and 3 and Lemma 4 of [19], a necessary and sufficient condition for exact recovery of sparse vectors via (7) is as follows. Let us define

$$\theta_{f_\sigma}(s, \mathbf{A}) \triangleq \sup_{\mathbf{h}\in\text{null}(\mathbf{A})\setminus\{\mathbf{0}\}} \frac{\sum_{i=1}^{s} f_\sigma(h_i^\downarrow)}{\sum_{i=1}^{m} f_\sigma(h_i)},$$

where $h_i^\downarrow$ denotes the ith largest (in magnitude) component of $\mathbf{h}$. $\theta_{f_\sigma}(s, \mathbf{A}) < 1/2$ is a necessary and sufficient condition for exact recovery of all vectors with sparsity at most s. These conditions are weaker than those corresponding to ℓ_1 minimization and lead to the following proposition which is a special case of [19, Proposition 5].

Proposition 3: Let $m^*_{f_\sigma}(\mathbf{A})$, $m^*_{\ell_0}(\mathbf{A})$, and $m^*_{\ell_1}(\mathbf{A})$ denote the maximum sparsity such that all vectors $\mathbf{x}$ with $\|\mathbf{x}\|_0 \le m^*_{f_\sigma}(\mathbf{A})$, $\|\mathbf{x}\|_0 \le m^*_{\ell_0}(\mathbf{A})$, and $\|\mathbf{x}\|_0 \le m^*_{\ell_1}(\mathbf{A})$ can be uniquely recovered by (7), (3), and (5), respectively. For any $f_\sigma(\cdot)$ possessing Property 1,

$$m^*_{\ell_1}(\mathbf{A}) \le m^*_{f_\sigma}(\mathbf{A}) \le m^*_{\ell_0}(\mathbf{A}).$$

The so-called robust recovery condition is satisfied if all vectors with sparsity at most s can be recovered from (8) with an error proportional to $\epsilon \ge \|\mathbf{w}\|$ [41]. In general, extension of the above necessary and sufficient conditions to robust recovery of sparse vectors from noisy measurement is not easy. However, [41] proves that, under some mild assumptions, the sets of sensing matrices satisfying the exact and robust recovery conditions differ by a set of measure zero. In other words, $\theta_{f_\sigma}(s, \mathbf{A}) < 1/2$ also guarantees the accurate recovery of sparse vectors via (8) for most sensing matrices.

Restricted isometry property based conditions: Let δ_{2s} and δ_{3s} denote the restricted isometry constants of orders $2s$ and $3s$ defined in [18]. For a general concave function $f_\sigma(\cdot)$ (including those satisfying Property 1), [40] shows that $\delta_{2s} < 1/2$ and $\delta_{3s} < 2/3$ are sufficient for accurate recovery of a sparse vector with the ℓ_0-norm not greater than s via (8).

Spherical section property based conditions: While the above recovery conditions do not provide strict superiority to ℓ_1 minimization, the following proposition shows that, in the noise-free case, one can obtain a solution arbitrarily close to the unique solution of ℓ_0 minimization by properly choosing σ. Indeed, as long as ℓ_0 minimization admits a unique solution, it is possible to recover it by (7) letting $\sigma \to 0$. However, it is obvious that, for sufficiently sparse vectors, $\theta_{f_\sigma}(s, \mathbf{A}) < 1/2$ guarantees exact recovery for any $\sigma > 0$. To state the result, first, the definition of the spherical section property is recalled.

Definition 1 (Spherical Section Property [20], [42]): The sensing matrix $\mathbf{A}$ possesses the Δ-spherical section property if, for all $\mathbf{w} \in \text{null}(\mathbf{A}) \setminus \{\mathbf{0}\}$, $\|\mathbf{w}\|_1^2/\|\mathbf{w}\|^2 \ge \Delta(\mathbf{A})$. In other words, the spherical section constant of the matrix $\mathbf{A}$ is defined as

$$\Delta(\mathbf{A}) \triangleq \min_{\mathbf{w}\in\text{null}(\mathbf{A})\setminus\{\mathbf{0}\}} \frac{\|\mathbf{w}\|_1^2}{\|\mathbf{w}\|^2}.$$

It is known that many randomly generated sensing matrices possess the SSP with high probability [42]. The proof of the following proposition easily follows from Proposition 4 of [22] by restricting the matrices to be diagonal.

Proposition 4: Assume that $\mathbf{A} \in \mathbb{R}^{n\times m}$ has the Δ-spherical property, and consider a class of functions $\{f_\sigma\}$ possessing Property 1. Let $\mathbf{x}_\sigma$ denote a solution to (7), and let $\mathbf{x}_0$ denote the unique solution to (3). Then

$$\|\mathbf{x}_\sigma - \mathbf{x}_0\| \leq \frac{n\alpha_\sigma}{\sqrt{\Delta} - \sqrt{\lceil \Delta - 1 \rceil}},$$

where $\alpha_\sigma = \left| f_\sigma^{-1}(1 - \frac{1}{n}) \right|$. The above inequality also leads to

$$\lim_{\sigma \to 0^+} \mathbf{x}_\sigma = \mathbf{x}_0.$$

V. Numerical Experiments

To assess the effectiveness of the SCSA algorithm in recovering sparse vectors, a number of numerical experiments are performed. Initially, the effect of parameter c is examined, and a suitable choice for this parameter is proposed. Next, the performance of SCSA in noiseless and noisy settings is evaluated and compared to some of the state-of-the-art algorithms. The general experimental setups as well as specific settings for the noise-free and noisy cases are described in the following subsection.

A. Experimental Setups

We use randomly generated sparse vectors and sensing matrices in all numerical experiments. More specifically, following common practice, each entry of the sensing matrix $\mathbf{A}$ is generated independently from the zero-mean, unit-variance Gaussian distribution $N(0,1)$, and the columns of $\mathbf{A}$ are normalized to have unit ℓ_2-norm. To construct a sparse vector $\widetilde{\mathbf{x}}$ with $\|\widetilde{\mathbf{x}}\|_0 = s$, first, the location of nonzero components is sampled uniformly at random among all possible subsets of $\{1, \cdots, m\}$ with cardinality s. Then the values of nonzero components are drawn independently from either $N(0,1)$ or the Rademacher distribution of $\{\pm 1\}$ with equal probability. Moreover, when measurements are noisy, the noise vector $\mathbf{w}$ is always drawn from $N(0, \sigma_w^2 \mathbf{I})$. Finally, the vector of measurements $\mathbf{b}$ is equal to $\mathbf{A}\widetilde{\mathbf{x}} + \mathbf{w}$, where $\mathbf{w} = \mathbf{0}$ when the noise-free case is under consideration.

Let $\widehat{\mathbf{x}}$ denote the output of one of the algorithms used in the numerical experiments to recover the sparse vector $\widetilde{\mathbf{x}}$ from either noisy or noiseless measurements. We use the following four quantities to measure and compare reconstruction accuracy in different experiments.

- Reconstruction SNR in dB:
 $\text{SNR}_{\text{rec}} \triangleq 20\log_{10}(\|\widetilde{\mathbf{x}}\|/\|\widetilde{\mathbf{x}} - \widehat{\mathbf{x}}\|)$ which will be used in Experiment 1 and, implicitly, in Experiment 2.
- Median reconstruction SNR:
 $\text{MSNR}_{\text{rec}} \triangleq 10\log_{10}(\|\widetilde{\mathbf{x}}\|^2/\operatorname{median}(\|\widetilde{\mathbf{x}} - \widehat{\mathbf{x}}\|^2))$ where $\operatorname{median}(\|\widetilde{\mathbf{x}} - \widehat{\mathbf{x}}\|^2)$ denotes the median of $\|\widetilde{\mathbf{x}} - \widehat{\mathbf{x}}\|^2$ over all the Monte-Carlo simulations. MSNR_{rec} is used in Experiments 3 and 4.
- Support recovery rate (SRR):
 Let $\widetilde{\tau}$ and $\widehat{\tau}$ denote the support set of $\widetilde{\mathbf{x}}$ and the set of indices of the s largest (in magnitude) components of $\widehat{\mathbf{x}}$, respectively. SRR is defined as the number of realizations in which $\widetilde{\tau} = \widehat{\tau}$ normalized by the total number of Monte-Carlo simulations. SRR is used in Experiment 4.
- Mean-squared error (MSE) which is the sample mean of $\|\widetilde{\mathbf{x}} - \widehat{\mathbf{x}}\|^2$ and will be used in Experiment 5.

Besides the accuracy, execution time, as a rough measure of the computational complexity, is used to compare algorithms. All simulations are performed in MATLAB 8 environment using an Intel Core i7-4600U, 2.1 GHz processor with 8 GB of RAM, under Microsoft Windows 7 operating system.

1) noise-free case: An algorithm is declared to be successful in recovering the solution, if $\text{SNR}_{\text{rec}} \geq 60$ dB. Consequently, to compare the performance of different algorithms, we use the success rate defined as the number of times an algorithm successfully recovers the solution divided by the total number of trials. To solve (5) and (11), where the latter is also used in the implementation of some of the algorithms in the comparison with an algorithm-dependent weighting matrix, we use the ℓ_1-magic [43].

2) noisy case: We assume that the noise variance, σ_w^2, is known; thus, it is possible to use the following formula [33], [44]

$$\lambda = 2c_r\sigma_w\Phi^{-1}(1 - \frac{\alpha_r}{2m}) \tag{24}$$

to choose the regularization parameter for the LASSO estimator (13). With the above choice, in which $c_r > 1$ is some constant and Φ is the cumulative density function of $N(0,1)$, the LASSO estimator achieves the so-called 'near-oracle' performance with probability at least $1 - \alpha_r$ [44]. [4] In our numerical experiments, c_r and α_r are set to 1.05 and 0.5, respectively. The above regularization parameter is used for IST and FISTA based implementations of the LASSO as well as IT and FIT instances of the SCSA algorithm.

Finally, it should be mentioned that, to have more stable plots without large fluctuations, in the noisy case, we normalize the ℓ_2 norm of each s-sparse vector to $\sqrt{s}$.

B. Effect of Parameter c

Experiment 1. To increase the accuracy of approximating the ℓ_0 norm in SCSA, σ is decreased according to the rule $\sigma_i = c\sigma_{i-1}, i \geq 1$. Intuitively, a small value for c corresponds to fast decay of σ and increase in the risk of getting trapped in local solutions. In contrast, a relatively large value of c leads to smooth changes in $F_\sigma(\cdot)$ where it is more likely to end up in the global solution. In the simplest case, where the sparsity of the solution is small enough and measurements are noise-free, the sufficient condition stated in Section IV implies that (5) and (7) have the same unique solution. As SCSA is initialized with the minimum ℓ_1-norm solution and Proposition 1 proves that, for any σ, $F_\sigma(\cdot)$, in the internal loop, is nonincreasing, the SCSA algorithm converges after 2 iterations independent of c. Moreover, in the noisy case, from the theoretical analysis presented in Section IV and the proof of Theorem 2 which shows that $\lambda_\sigma F_\sigma(|\mathbf{x}_i|) + h(\mathbf{x}_i)$ is not increasing in i, we expect a similar behavior. On the other hand, when the sparsity level increases, to decrease the risk of getting trapped in local minima, a larger c should be selected.

[4] As will be explained later in Subsection V-D, with this choice of α_r which is the same as in [45], we are evaluating the typical performance of the LASSO, SCSA-IT, and SCSA-FIT algorithms.

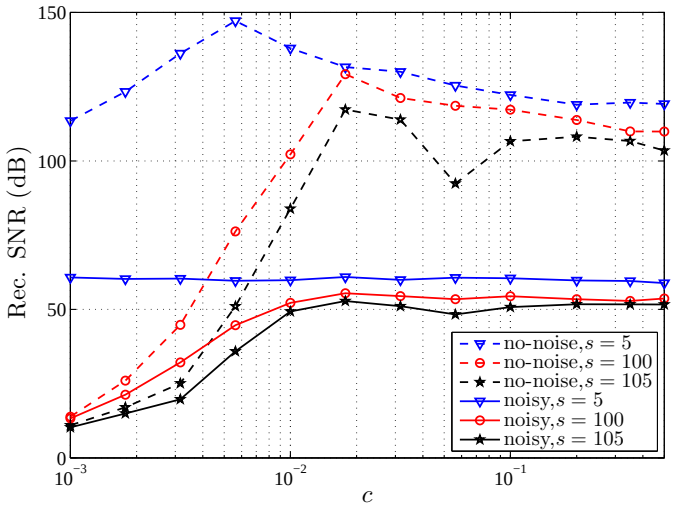


Fig. 3. Averaged SNR_{rec} of the SCSA algorithm in solving noiseless and noisy problems are plotted versus c for 3 different sparsity levels. The sparse vector is of length 400, and the sensing matrix is 200×400. To have an accurate estimate of the SNR_{rec}, in each problem, 100 Monte-Carlo simulations are run, and results are averaged.

To see the above intuition, the effect of parameter c in the reconstruction SNR in noisy and noiseless cases is numerically experimented. The dimensions of the sensing matrix are fixed to 200×400, σ_w is equal to 10^{-3}, and the sparse vectors are Gaussian distributed. The experiment is repeated for 3 different cardinalities ($s = 5, 100, 105$), and SNR_{rec} is averaged over 100 trials. Fig. 3 shows the averaged SNR_{rec}'s as a function of c. As predicted, when s is small, the SNR_{rec} is always high. On the other hand, for high sparsity, after passing a critical value, SNR_{rec} remains almost unchanged. Based on this observation, in the remaining experiments, we conservatively set c to 0.1.

C. Noise-Free Recovery

Experiment 2. In this experiment, the performance of the SCSA algorithm in the noise-free setting is compared to ℓ_1 minimization, the SL0 algorithm [14], ℓ_p quasi-norm minimization [13], [15], and the reweighted ℓ_1 minimization [12] in terms of success rate and execution time. The following implementation method and parameters are used for each algorithm.

- When implementing the SCSA algorithm, ϵ_1 and ϵ_2 are set to 10^{-3} and 10^{-2}, respectively, and c is set to 0.1.
- For SL0 (MATLAB code: http://ee.sharif.edu/~SLzero/), the following parameters are used: `sigma_min=`10^{-4}, `c=0.8`, `mu=2`, and `L=8`. As suggested in [14], these parameters result in a much better success rate than that provided by the default values.
- ℓ_p quasi-norm minimization is implemented based on the iteratively reweighted ℓ_1 minimization approach in [46] with $p = 0.5$.
- For the reweighted ℓ_1 minimization [12], we use $\epsilon = 0.1$ which attains the best performance [12].

The dimensions of the sensing matrix are fixed to 250×500, the sparsity of $\widetilde{\mathbf{x}}$ is changed from 70 to 170, and success rate and average execution time over 500 trials are plotted in

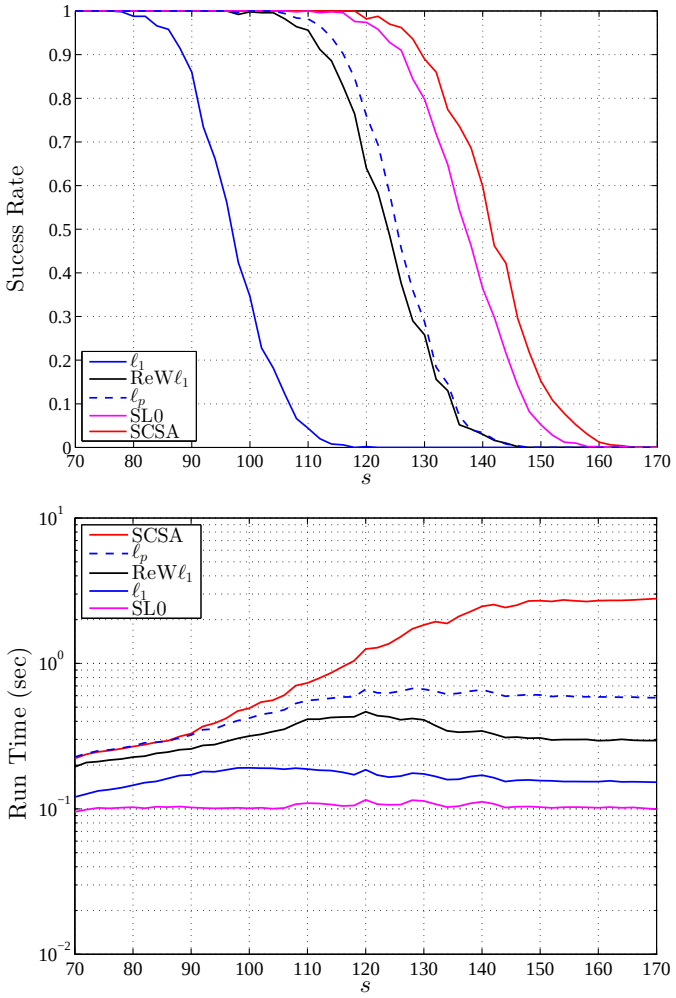


Fig. 4. Comparison of the SCSA algorithm in solving the noise-free problem to ℓ_1 minimization, the SL0 algorithm [14], ℓ_p quasi-norm minimization [13], [15] ($p = 0.5$), and the reweighted ℓ_1 minimization [12] in terms of success rate and execution time. The dimensions of the sensing matrix are 250×500. Trials are repeated 500 times, and results are averaged over them.

Fig. 4. As depicted in this figure, SCSA has the best success rate amongst the algorithms, whereas its computational load is much higher than the closest competitor. However, as we show shortly, in the noisy setting which is more realistic, SCSA maintains the superiority with a quite reasonable complexity.

D. Recovery from Noisy Measurements

In the following experiments, superiority of the proposed algorithm in the noisy setting is demonstrated. Toward this end, the SCSA algorithm is compared with the oracle estimator [21], which knows the location of the nonzero elements of the true solution, LASSO (or BPDN), the SCAD penalty [47], Robust SL0 [48], the method of [49], [50], and iterative log thresholding (ILT) [51]. The description and implementation details of these algorithm are as follows.

To implement the LASSO estimator, IST and/or FISTA methods are used. The smoothly clipped absolute deviation (SCAD) penalty is a well-known nonconvex function that promotes sparsity more tightly than the ℓ_1 norm does and has some oracular properties [47]. For this penalty, we set the

parameter a to 3.7 as suggested in [47]. To efficiently solve the optimization problem resulting from the SCAD penalty, we use the algorithm of [52] with parameter $\tau = 10^{-3}$ and the external loop stopping threshold equal to 10^{-4}. [5] Robust SL0 (RSL0), a modification to the original SL0 to handle noisy measurements, instead of a regularization parameter, needs the noise variance in the denoising step. To have a fair comparison, σ_w is passed to it. Other parameters of RSL0 are the same as in the Experiment 2 except for `sigma_min`=$\sigma_w/10$. The method of [49], [50], which we refer to as IST-p, uses a generalized version of the soft thresholding operator, $\mathcal{S}_\alpha$, defined as

$$\mathcal{S}_\alpha^{(p)}(x) = \max(|x| - \alpha|x|^{p-1}, 0)\,\mathrm{sign}(x);$$

otherwise, it is identical to the IST. We use two instances of this method with $p = 0.5$ and $p = 0.1$. The ILT algorithm is an extension of the reweighted ℓ_1 minimization in [12] to the noisy case. In fact, it solves the following optimization problem

$$\min_{\mathbf{x}} \lambda \sum_{i=1}^{m} \log(|x_i| + \beta) + \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2,$$

in which β is some small constant to ensure positivity of the argument of $\log(\cdot)$, using iterative thresholding approach.

For SCSA-FIT, SCSA-IT, IST, and FISTA, the regularization parameter is chosen according to the formula given in (24). For ILT, SCAD, and IST-p, the regularization parameter is numerically tuned at $\sigma_w = 10^{-2}$ and is linearly scaled with the change of noise standard deviation. The stopping criterion for IST, FISTA, IST-p, and ILT is $d = \|\mathbf{x}_i - \mathbf{x}_{i-1}\|/\|\mathbf{x}_{i-1}\| \le \eta$, where $\mathbf{x}_i$ and $\mathbf{x}_{i-1}$ are the solutions at the ith and $(i-1)$th iterations, $\eta = \min(10^{-3}\lambda, 10^{-4})$, and λ is the associated regularization parameter. For SCSA-FIT and SCSA-IT, ϵ_1 and ϵ_2 are set as suggested in Remark 2. Moreover, for IST, FISTA, ILT, and IST-p, μ is always fixed to $0.99/(2\lambda_{max}(\mathbf{A}^T\mathbf{A}))$, while for SCSA-FIT and SCSA-IT, μ is $0.99/(2\lambda_{max}(\mathbf{A}^T\mathbf{A}) + \lambda/\sigma)$.

To obtain accurate and stable results, similar to [45], $\mathrm{MSNR}_{\mathrm{rec}}$ is used in Experiments 3 and 4 to compare the performance of the algorithms. In fact, with $\mathrm{MSNR}_{\mathrm{rec}}$, we are comparing the 'typical' performance of these algorithms as most performance guarantees for the LASSO estimator and other nonconvex estimators hold with some probability (see e.g., [26], [33], [45]).

Experiment 3. Under the above conditions, for Gaussian distributed sparse vectors, s is changed from 2 to 160, and the $\mathrm{MSNR}_{\mathrm{rec}}$ and the averaged execution time (except for the oracle estimator) for all the algorithms over 500 runs are plotted in Fig. 5. As clearly demonstrated in this figure, the (median) reconstruction SNR of the SCSA-FIT algorithm, for almost all values of s, is higher than others. Also, it has a near-oracle performance for a broader range of sparsity levels. So far as the computational load is concerned, SCSA-FIT needs at most (approximately) 3 times higher execution time in comparison to FISTA, the fastest algorithm for most of sparsity levels. However, the computational cost is lower than

[5]MATLAB code: https://sites.google.com/site/alainrakotomamonjy/

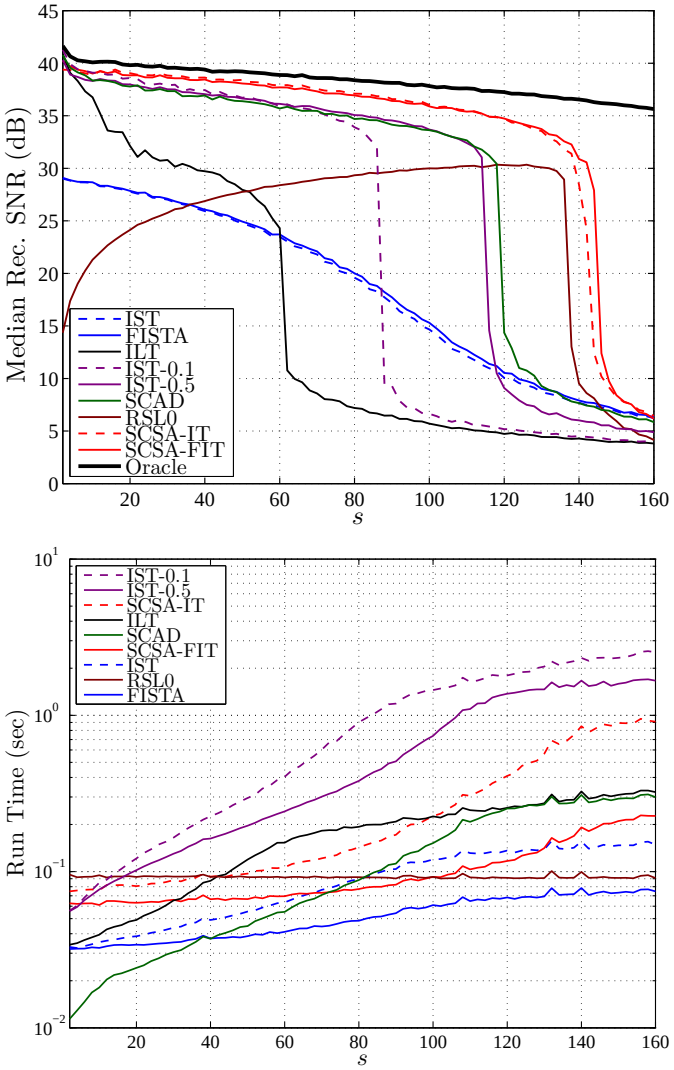


Fig. 5. Comparison of two variants of the SCSA algorithm in solving the noisy problem to IST, FISTA, SCAD [47], ILT [51], RSL0 [48], IST-p [49], [50], and the oracle estimator [21] in terms of $\mathrm{MSNR}_{\mathrm{rec}}$ and execution time. The dimensions of the sensing matrix are 250×500. The sparse vectors are Gaussian distributed, and $\sigma_w = 10^{-2}$. Trials are repeated 500 times, and results are averaged over them.

that of SCAD and RSL0, when s is larger than 70 and is smaller than 100, respectively. These two algorithms (SCAD and RSL0) are somehow the best competitors; however, they are not able to follow the oracular performance at the sparsity level that SCSA does.

SCSA-IT and SCSA-FIT have a quite similar performance in terms of $\mathrm{MSNR}_{\mathrm{rec}}$. However, as expected, the former spends considerably more time than the latter to output a solution. Particularly, SCSA-FIT is approximately 8 times faster than SCSA-IT, when sparsity level is equal to 140.

Experiment 4. To show that the SCSA algorithm is effective for sparse vectors which are not Gaussian-like distributed, we repeat Experiment 3 with the Rademacher distributed signals. The Rademacher distributed sparse vectors do not exhibit the power-law decaying behavior [43] when nonzero components are sorted according to their magnitude. In addition, some numerical simulations show that SL0, which

parallels the idea of SCSA, does not work well for this kind of distributions (see nuit-blanche.blogspot.com/2011/11/post-peer-review-of-sl0.html).

In many applications, it is more crucial to find the support set of the sparse vector accurately [53]. To numerically asses the performance of SCSA in recovering the true support, we also calculate the SRR in this experiment. Fig. 6 illustrates the results of the experiment for SCSA as well as all other algorithms in Experiment 3. As shown in this figure, SCSA and some other algorithms follow the oracle performance more closely than they did in Experiment 3. Furthermore, SCSA-FIT achieves the best performance in terms of MSNR_{rec} and SRR, whereas its complexity is quite comparable to FISTA.

Experiment 5. In this experiment, we examine and compare the accuracy of the SCSA-FIT algorithm in terms of MSE to FISTA, ILT, IST-0.5, SL0, and the oracle estimator as the noise variance changes. Under the same conditions as in Experiment 3 and for 3 sparsity levels ($s = 10, 50, 105$), the noise standard deviation is changed from 10^{-1} to 10^{-3}, and the MSE is calculated. Results of this experiment are summarized in Fig. 7. As clearly depicted, for all values of s, SCSA-FIT is the closest one to the oracle estimator, and can follow it even for the large sparsity level of $s = 105$.

VI. Conclusion

For the problem of recovering sparse vectors from compressed measurements, we proposed to replace the ℓ_0 norm with a family of concave functions that closely approximates sparsity. To solve the consequent nonconvex optimization problem, we exploited a continuation approach leading to starting from ℓ_1 minimization, followed by successively solving accuracy-increasing approximations of the ℓ_0 norm subject to the constraints. In the presence of noise in the measurements, we combined the continuation approach to iterative thresholding method of optimization and proposed a computationally inexpensive algorithm. This was obtained by deriving the closed-form solution for the program (18). We also provided a choice for the regularization parameter in (16) which is based on the regularization parameter for the LASSO estimator. The numerical simulations revealed that the proposed method considerably and consistently outperforms some of the common algorithms (including LASSO), especially when the sparsity level increases.

Appendix A

Before deriving the closed-form solution to (18) for $f_\sigma(|x|) = 1-\exp(-|x|/\sigma)$, we need to introduce the Lambert W function and prove an inequality on this function.

The Lambert W function, which has several applications in physics and applied and pure mathematics (particularly, combinatorics) [35], is defined as the multivalued inverse of the function $w \mapsto we^w$ [35]. More precisely, the Lambert W function, denoted by $W(x)$, is given implicitly by

$$W(x)e^{W(x)} = x$$

where x in general can be a complex number, but, herein, we only deal with real-valued W and assume x to be real. In this fashion, $W(x)$ is single-valued for $x \geq 0$ and is double-valued for $-\frac{1}{e} \leq x < 0$ [35]. To discriminate between the two branches for $x \in [-\frac{1}{e}, 0)$, we use the same notation as in [35] and denote the branch satisfying $W(x) \geq -1$ and $W(x) \leq -1$ by $W_0(x)$ and $W_{-1}(x)$, respectively.

Lemma 1: For any $y \in [-\frac{1}{e}, 0)$, $W_0(y) + W_{-1}(y) \leq -2$.

Proof: Simple algebraic manipulation shows that the parabola $e^{-1}(\frac{1}{2}x^2 + x - \frac{1}{2})$ is below the function xe^x for $x \in (-1, 0)$ and above it for $x < -1$. Therefore, any line parallel to the x-axis with the y-intercept in the interval $y = [-\frac{1}{e}, 0)$, intersects the parabola at equal or larger x coordinates than those corresponding to intersections with xe^x. From the other side, the x coordinate of the intersection points of these lines with xe^x gives $W_0(y)$ and $W_{-1}(y)$. Consequently, for any $y \in [-\frac{1}{e}, 0)$, the sum $W_0(y) + W_{-1}(y)$ is less than or equal to the sum of the roots of $e^{-1}(\frac{1}{2}x^2 + x - \frac{1}{2}) = y$ which is always equal to -2. This completes the proof. ∎

To obtain the solution to (18), we first notice that since $F_\sigma(\cdot)$ is a separable function, the minimizer can be obtained element by element. Therefore, we focus on the one-variable optimization problem. To that end, let $L(x, x_0) = \frac{1}{2\mu}(x - x_0)^2 + \lambda f_\sigma(|x|)$ denote the corresponding scalar version of (18) and let $x^* = \operatorname{argmin}_x L(x, x_0)$. It is easy to show that $L(x, -x_0) = L(-x, x_0)$; thus, from

$$\begin{cases} x^* = \operatorname{argmin}_x\{L(x, x_0)\} & \text{for } x_0 \geq 0 \\ x^* = -\operatorname{argmin}_x\{L(-x, x_0)\} & \text{for } x_0 < 0, \end{cases}$$

we get

$$\begin{aligned} x^* &= \operatorname{sign}(x_0) \operatorname*{argmin}_x \Big\{ \frac{1}{2\mu}(\operatorname{sign}(x_0)x - x_0)^2 + \lambda f_\sigma(|x|) \Big\}, \\ &= \operatorname{sign}(x_0) \operatorname*{argmin}_x \Big\{ \frac{1}{2\mu}(x - |x_0|)^2 + \lambda f_\sigma(|x|) \Big\}. \end{aligned}$$

Let $y^* = \operatorname{argmin}_x\{\frac{1}{2\mu}(x - |x_0|)^2 + \lambda f_\sigma(|x|)\}$. This definition implies that $y^* \geq 0$, since if $y^* < 0$, then $\widehat{y} = -y^*$ contradicts the optimality of y^* as $(\widehat{y} - |x_0|)^2 < (y^* - |x_0|)^2$. Consequently, we have

$$x^* = \operatorname{sign}(x_0) \operatorname*{argmin}_x \Big\{ \frac{1}{2\mu}(x - |x_0|)^2 + \lambda f_\sigma(x) \mid x \geq 0 \Big\}.$$

Denoting $\widehat{L}(x, x_0) = \frac{1}{2\mu}(x - |x_0|)^2 + \lambda f_\sigma(x)$, putting $f_\sigma(x) = 1 - \exp(-x/\sigma)$, and differentiating $\widehat{L}(x, x_0)$ with respect to x, it can be obtained

$$\frac{1}{\mu}(x - |x_0|) = -\frac{\lambda}{\sigma}e^{-\frac{x}{\sigma}} \Longrightarrow \frac{x - |x_0|}{\sigma}e^{\frac{x-|x_0|}{\sigma}} = -\frac{\mu\lambda}{\sigma^2}e^{-\frac{|x_0|}{\sigma}}.$$

The above equation admits two solutions

$$\begin{cases} x_1 = \sigma W_0(-\frac{\mu\lambda}{\sigma^2}e^{-\frac{|x_0|}{\sigma}}) + |x_0| \\ x_2 = \sigma W_{-1}(-\frac{\mu\lambda}{\sigma^2}e^{-\frac{|x_0|}{\sigma}}) + |x_0|. \end{cases}$$

Since, for $x < -\frac{1}{e}$, $W(x)$ is not defined, we should have $-\frac{\mu\lambda}{\sigma^2}e^{-\frac{|x_0|}{\sigma}} \geq -\frac{1}{e}$ or $|x_0| \geq \sigma(1 + \ln(\mu\lambda) - 2\ln(\sigma))$; otherwise, the minimizer of $\widehat{L}(x, x_0)$ lies at $x = 0$.

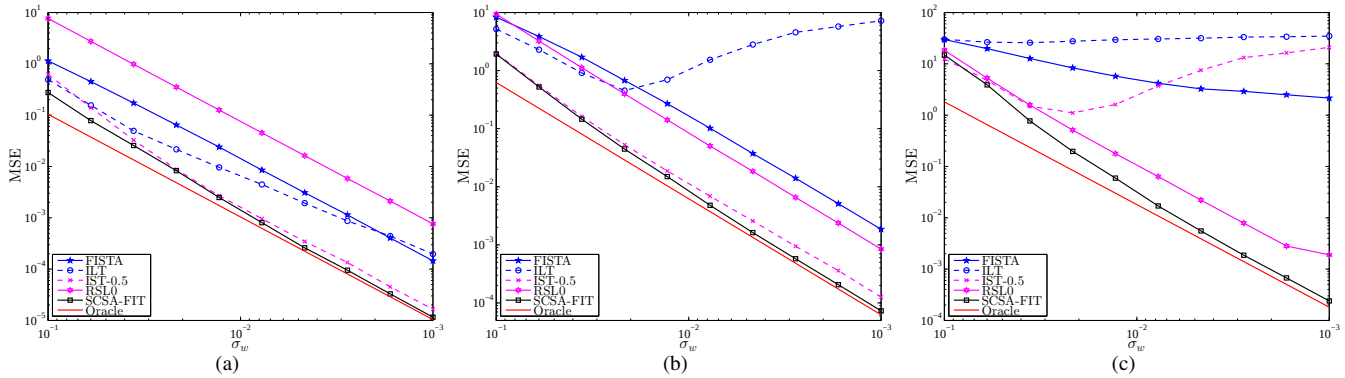


Fig. 7. The mean-squared errors of SCSA-FIT, FISTA, ILT, IST-0.5, RSL0, and the oracle estimator in recovering Gaussian distributed sparse vectors from noisy measurements as a function of the noise standard deviation. The dimensions of the sensing matrix are 250×500. Trials are repeated 500 times, and results are averaged over them. (a), (b), and (c) correspond to $s = 10, 50, 105$, respectively.

Substituting x_1 and x_2 in $\widehat{L}(x, x_0)$, letting $z = -\frac{\mu\lambda}{\sigma^2} e^{-\frac{|x_0|}{\sigma}}$, and using the definition $W(x)e^{W(x)} = x$, we have

$$\widehat{L}(x_1, x_0) = \lambda + \frac{\sigma^2}{\mu}\Big(\frac{1}{2}W_0^2(z) + W_0(z)\Big),$$
$$\widehat{L}(x_2, x_0) = \lambda + \frac{\sigma^2}{\mu}\Big(\frac{1}{2}W_{-1}^2(z) + W_{-1}(z)\Big).$$

From Lemma 1, we have

$$\begin{cases} W_0(z) + W_{-1}(z) \le -2 \\ W_0(z) \ge -1 \end{cases}$$
$$\Rightarrow 0 \le W_0(z) + 1 \le -1 - W_{-1}(z)$$
$$\Rightarrow \frac{1}{2}W_0^2(z) + W_0(z) \le \frac{1}{2}W_{-1}^2(z) + W_{-1}(z).$$

This shows that x_2 cannot be a minimizer of $\widehat{L}(x, x_0)$. In addition, it can be easily checked that $\widehat{L}(x, x_0)$ for $x \ge 0$ is convex only when $\sigma^2 \ge \mu\lambda$. Therefore, it can occur that the minimizer resides at the border ($x = 0$). It is very hard to analytically find the condition under which $x = x_1$ or $x = 0$ is the minimizer. However, one can easily compare the cost function at these points to find the minimizer. In summary, the one dimensional shrinkage operator can be characterized as

$$\mathcal{T}_{\mu\lambda}^{(\sigma)}(x_0) = \begin{cases} 0 & |x_0| \ge \sigma(1 + \ln(\mu\lambda/\sigma^2)) \\ 0 & \widehat{L}(x_1, x_0) \ge \widehat{L}(0, x_0) \\ \sigma W_0(z) + |x_0| & \text{otherwise,} \end{cases} \tag{25}$$

where $z = -\frac{\mu\lambda}{\sigma^2} e^{-\frac{|x_0|}{\sigma}}$.

Appendix B
Proof of Theorem 2

For the sake of simplicity, let us define $\phi(\mathbf{x}) \triangleq \lambda_\sigma F_\sigma(\mathbf{x})$, where, without introducing any ambiguity, we omitted the subscript σ. Further, let $g(\mathbf{x}) \triangleq h(\mathbf{x}) + \phi(|\mathbf{x}|)$ denote the cost function in (16). The program (17) now is equal to

$$\mathbf{x}_{k+1} = \underset{\mathbf{x}}{\operatorname{argmin}} \Big\{ \langle \mathbf{x} - \mathbf{x}_k, \nabla h(\mathbf{x}_k)\rangle + \frac{1}{2\mu}\|\mathbf{x} - \mathbf{x}_k\|^2 + \phi(|\mathbf{x}|) \Big\}. \tag{26}$$

Since $h(\mathbf{x})$ has an M-Lipschitz continuous gradient, we have [54]

$$h(\mathbf{x}_{k+1}) \le h(\mathbf{x}_k) + \langle \mathbf{x}_{k+1} - \mathbf{x}_k, \nabla h(\mathbf{x}_k)\rangle + \frac{M}{2}\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2. \tag{27}$$

Moreover, $\nabla^2\phi(\mathbf{x}) \succeq -M'\mathbf{I}$ for $\mathbf{x} \ge 0$, implies that, for all $\mathbf{x}, \mathbf{y} \ge \mathbf{0}$,

$$\phi(\mathbf{x}) \le \phi(\mathbf{y}) + \langle \mathbf{x} - \mathbf{y}, \nabla\phi(\mathbf{x})\rangle + \frac{M'}{2}\|\mathbf{y} - \mathbf{x}\|^2.$$

Substituting $\mathbf{x}$ and $\mathbf{y}$ with $|\mathbf{x}_{k+1}|$ and $|\mathbf{x}_k|$, respectively, we get

$$\phi(|\mathbf{x}_{k+1}|) \le \phi(|\mathbf{x}_k|) + \langle |\mathbf{x}_{k+1}| - |\mathbf{x}_k|, \nabla\phi(|\mathbf{x}_{k+1}|)\rangle + \frac{M'}{2}\||\mathbf{x}_{k+1}| - |\mathbf{x}_k|\|^2.$$

Using $\||\mathbf{x}| - |\mathbf{y}|\| \le \|\mathbf{x} - \mathbf{y}\|$, the above inequality resorts to

$$\phi(|\mathbf{x}_{k+1}|) \le \phi(|\mathbf{x}_k|) + \langle |\mathbf{x}_{k+1}| - |\mathbf{x}_k|, \nabla\phi(|\mathbf{x}_{k+1}|)\rangle + \frac{M'}{2}\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2. \tag{28}$$

From the other side, $\mathbf{x}_{k+1}$ is a minimizer of (26), so we can write

$$\langle \mathbf{x}_{k+1} - \mathbf{x}_k, \nabla h(\mathbf{x}_k)\rangle \le \phi(|\mathbf{x}_k|) - \phi(|\mathbf{x}_{k+1}|) - \frac{1}{2\mu}\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2. \tag{29}$$

First-order concavity condition for ϕ implies that $\phi(\mathbf{y}) \le \phi(\mathbf{x}) + \langle \mathbf{y} - \mathbf{x}, \nabla\phi(\mathbf{x})\rangle$. Replacing $\mathbf{x}$ and $\mathbf{y}$ with $|\mathbf{x}_{k+1}|$ and $|\mathbf{x}_k|$, respectively, we get

$$\phi(|\mathbf{x}_k|) - \phi(|\mathbf{x}_{k+1}|) + \langle |\mathbf{x}_{k+1}| - |\mathbf{x}_k|, \nabla\phi(|\mathbf{x}_{k+1}|)\rangle \le 0. \tag{30}$$

Combining (27) and (28) results in

$$\begin{aligned} g(\mathbf{x}_{k+1}) - g(\mathbf{x}_k) &\le \langle \mathbf{x}_{k+1} - \mathbf{x}_k, \nabla h(\mathbf{x}_k)\rangle + \langle |\mathbf{x}_{k+1}| - |\mathbf{x}_k|, \nabla\phi(|\mathbf{x}_{k+1}|)\rangle \\ &\quad + \frac{1}{2}(M + M')\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 \\ &\overset{(a)}{\le} \phi(|\mathbf{x}_k|) - \phi(|\mathbf{x}_{k+1}|) + \langle |\mathbf{x}_{k+1}| - |\mathbf{x}_k|, \nabla\phi(|\mathbf{x}_{k+1}|)\rangle \\ &\quad + \frac{1}{2}(M + M' - \frac{1}{\mu})\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2, \end{aligned} \tag{31}$$

where (a) follows from (29). Assume that $0 < \mu < \frac{1}{M+M'}$, then (31) rearranges to

$$\begin{aligned}\frac{1}{2}(\frac{1}{\mu} - M - M')\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 \\ &\leq g(\mathbf{x}_k) - g(\mathbf{x}_{k+1}) + \phi(|\mathbf{x}_k|) - \phi(|\mathbf{x}_{k+1}|) \\ &\quad + \langle |\mathbf{x}_{k+1}| - |\mathbf{x}_k|, \nabla\phi(|\mathbf{x}_{k+1}|)\rangle \\ &\overset{(b)}{\leq} g(\mathbf{x}_k) - g(\mathbf{x}_{k+1}), \end{aligned} \tag{32}$$

where (b) follows from (30).

Now, we show that $\{g(\mathbf{x}_k)\}$ is a nonincreasing sequence. Similar to (27), one has

$$h(\mathbf{x}) \leq h(\mathbf{x}_k) + \langle \mathbf{x} - \mathbf{x}_k, \nabla h(\mathbf{x}_k)\rangle + \frac{M}{2}\|\mathbf{x} - \mathbf{x}_k\|^2.$$

As a result, if $\mu \in (0, \frac{1}{M})$, then $h(\mathbf{x}) \leq H(\mathbf{x}, \mathbf{x}_k) \triangleq h(\mathbf{x}_k) + \langle \mathbf{x} - \mathbf{x}_k, \nabla h(\mathbf{x}_k)\rangle + \frac{1}{2\mu}\|\mathbf{x} - \mathbf{x}_k\|^2$, or, equivalently,

$$g(\mathbf{x}) \leq H(\mathbf{x}, \mathbf{x}_k) + \phi(|\mathbf{x}|) \quad \forall \mathbf{x}.$$

Replacing $\mathbf{x}$ with $\mathbf{x}_{k+1}$, we get $g(\mathbf{x}_{k+1}) \leq H(\mathbf{x}_{k+1}, \mathbf{x}_k) + \phi(|\mathbf{x}_{k+1}|)$. Furthermore, $\mathbf{x}_{k+1}$ is a solution to (26), so $H(\mathbf{x}_{k+1}, \mathbf{x}_k) + \phi(|\mathbf{x}_{k+1}|) \leq H(\mathbf{x}_k, \mathbf{x}_k) + \phi(|\mathbf{x}_k|) = g(\mathbf{x}_k)$. This inequality together with the previous one leads to

$$g(\mathbf{x}_{k+1}) \leq g(\mathbf{x}_k) \quad \forall k \geq 0.$$

Summing (32) over $k = 0, \cdots, N$, we get

$$\frac{1}{2}(\frac{1}{\mu} - M - M')\sum_{k=0}^{N}\|\mathbf{x}_{k+1} - \mathbf{x}_k\|^2 \leq g(\mathbf{x}_0) - g(\mathbf{x}_{N+1}).$$

Since $g(\mathbf{x}_0)$ is finite, we conclude that $\{\mathbf{x}_k\}$ is convergent. Next, we prove that $\{\mathbf{x}_k\}$ converges to a stationary point of (16). Suppose that $\{\mathbf{x}_k\}$ converges to $\mathbf{x}^*$. It can be verified that the cost function in (26) is strictly convex for $\mu < 1/M$; thus, the minimizer of (26) is unique. This fact together with [55, Thm. 1.17 and Thm. 7.41] implies that when $k \to \infty$, we have

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \Big\{ \langle \mathbf{x} - \mathbf{x}^*, \nabla h(\mathbf{x}^*)\rangle + \frac{1}{2\mu}\|\mathbf{x} - \mathbf{x}^*\|^2 + \phi(|\mathbf{x}|) \Big\}.$$

First-order optimality condition of the above program leads to

$$\mathbf{0} \in \nabla h(\mathbf{x}^*) + \frac{1}{\mu}(\mathbf{x} - \mathbf{x}^*) + \partial\phi(|\mathbf{x}|) \quad \text{at } \mathbf{x} = \mathbf{x}^*,$$

proving that $\mathbf{x}^*$ is a stationary point of (16). ■

REFERENCES

[1] D. Gross, Y. K. Liu, S. T. Flammia, S. Becker, and J. Eisert, "Quantum state tomography via compressed sensing," *Physical review letters*, vol. 105, no. 15, pp. 150401, 2010.

[2] Y. Wiaux, L. Jacques, G. Puy, A. Scaife, and P. Vandergheynst, "Compressed sensing imaging techniques for radio interferometry," *Mon. Not. R. Astron. Soc.*, vol. 395, no. 3, pp. 1733–1742, 2009.

[3] D. Koslicki, S. Foucart, and G. Rosen, "Quikr: a method for rapid reconstruction of bacterial communities via compressive sensing," *Bioinformatics*, vol. 29, no. 17, pp. 2096–2102, 2013.

[4] Y. Erlich, A. Gordon, M. Brand, G. Hannon, and P. Mitra, "Compressed genotyping," *IEEE Trans. Inf. Theory*, vol. 56, no. 2, pp. 706–723, 2010.

[5] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic resonance in medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.

[6] Y. Eldar, "Compressed sensing of analog signals in shift-invariant spaces," *IEEE Trans. Signal Process.*, vol. 57, no. 8, pp. 2986–2997, 2009.

[7] I Bilik, "Spatial compressive sensing for direction-of-arrival estimation of multiple sources using dynamic sensor arrays," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 47, no. 3, pp. 1754–1769, 2011.

[8] M. Herman and T. Strohmer, "High-resolution radar via compressed sensing," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2275–2284, 2009.

[9] D. Donoho, M. Elad, and V. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory*, vol. 52, no. 1, pp. 6–18, 2006.

[10] M. Babaie-Zadeh and C. Jutten, "On the stable recovery of the sparsest overcomplete representations in presence of noise," *IEEE Trans. Signal Process.*, vol. 58, no. 10, pp. 5396–5400, 2010.

[11] M. Elad, *Sparse and redundant representations: from theory to applications in signal and image processing*, Springer, 2010.

[12] E.J. Candés, M.B. Wakin, and S.P. Boyd, "Enhancing sparsity by reweighted ℓ_1 minimization," *Journal of Fourier analysis and applications*, vol. 14, no. 5-6, pp. 877–905, 2008.

[13] R. Chartrand, "Exact reconstruction of sparse signals via nonconvex minimization," *IEEE Signal Process. Lett.*, vol. 14, no. 10, pp. 707–710, 2007.

[14] H. Mohimani, M. Babaie-Zadeh, and C. Jutten, "A fast approach for overcomplete sparse decomposition based on smoothed ℓ_0 norm," *IEEE Trans. Signal Process.*, vol. 57, no. 1, pp. 289–301, 2009.

[15] S. Foucart and M. Lai, "Sparsest solutions of underdetermined linear systems via ℓ_q-minimization for $0 < q \leq 1$," *Appl. Comput. Harmon. Anal.*, vol. 26, no. 3, pp. 395–407, 2009.

[16] S. Rangan, "Generalized approximate message passing for estimation with random linear mixing," in *Proc. IEEE Int. Symp. Inform. Theory*, 2011, pp. 2168–2172.

[17] P.L. Combettes and J. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-point algorithms for inverse problems in science and engineering*, pp. 185–212. Springer, 2011.

[18] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Comptes Rendus Mathematique*, vol. 346, no. 9, pp. 589–592, 2008.

[19] R. Gribonval and M. Nielsen, "Highly sparse representations from dictionaries are unique and independent of the sparseness measure," *Appl. Comput. Harmon. Anal.*, vol. 22, pp. 335–355, 2007.

[20] B.S. Kashin and V.N. Temlyakov, "A remark on compressed sensing," *Mathematical notes*, vol. 82, no. 5-6, pp. 748–755, 2007.

[21] E. Candés and T. Tao, "The dantzig selector: Statistical estimation when p is much larger than n," *The Annals of Statistics*, pp. 2313–2351, 2007.

[22] M. Malek-Mohammadi, M. Babaie-Zadeh, and M. Skoglund, "Iterative concave rank approximation for recovering low-rank matrices," *IEEE Trans. Signal Process.*, vol. 62, no. 20, pp. 5213–5226, 2014.

[23] J. Trzasko and A. Manduca, "Highly undersampled magnetic resonance image reconstruction via homotopic-minimization," *IEEE Transactions on Medical imaging*, vol. 28, no. 1, pp. 106–121, 2009.

[24] R. Chartrand and V. Staneva, "Restricted isometry properties and nonconvex compressive sensing," *Inverse Problems*, vol. 24, no. 3, 2008.

[25] R. Wu and D. Chen, "The improved bounds of restricted isometry constant for recovery via ℓ_p minimization," *IEEE Trans. Inf. Theory*, vol. 59, 2013.

[26] T. Zhang, "Analysis of multi-stage convex relaxation for sparse regularization," *J. Mach. Learn. Res.*, vol. 11, pp. 1081–1107, 2010.

[27] C. Zhang and T. Zhang, "A general theory of concave regularization for high-dimensional sparse estimation problems," *Statistical Science*, vol. 27, no. 4, pp. 576–593, 2012.

[28] R. Chartrand, "Nonconvex compressed sensing and error correction," in *IEEE Int. Conf. Acoust. Speech Signal Process.*, 2007, pp. 889–892.

[29] Y. Shen, J. Fang, and H. Li, "Exact reconstruction analysis of log-sum minimization for compressed sensing," *IEEE Signal Process. Lett.*, vol. 20, no. 12, pp. 1223–1226, 2013.

[30] D. Needell, "Noisy signal recovery via iterative reweighted l1-minimization," in *Asilomar Conference on Signals, Systems, and Computers*, 2009, pp. 113–117.

[31] M. Malek-Mohammadi, M. Babaie-Zadeh, A. Amini, and C. Jutten, "Recovery of low-rank matrices under affine constraints via a smoothed rank function," *IEEE Trans. Signal Process.*, vol. 62, no. 4, pp. 981–992, 2014.

[32] D. Hunter and K. Lange, "A tutorial on MM algorithms," *The American Statistician*, vol. 58, no. 1, pp. 30–37, 2004.

[33] P. Bickel, Y. Ritov, and A. Tsybakov, "Simultaneous analysis of lasso and dantzig selector," *The Annals of Statistics*, pp. 1705–1732, 2009.

[34] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.

[35] R. M. Corless, G. H. Gonnet, D. EG Hare, D. J. Jeffrey, and D. E. Knuth, "On the Lambert W function," *Advances in Computational mathematics*, vol. 5, no. 1, pp. 329–359, 1996.

[36] F. Clarke, *Optimization and nonsmooth analysis*, SIAM, Philadelphia, 1990.

[37] T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *Appl. Comput. Harmon. Anal.*, vol. 27, no. 3, pp. 265–274, 2009.

[38] T. Zhang, "Multi-stage convex relaxation for feature selection," *Bernoulli*, vol. 19, no. 5B, pp. 2277–2293, 2013.

[39] Zh. Wang, H. Liu, and T. Zhang, "Optimal computational and statistical rates of convergence for sparse nonconvex learning problems," *Annals of statistics*, vol. 42, no. 6, pp. 2164, 2014.

[40] C. Zhang, "Nearly unbiased variable selection under minimax concave penalty," *The Annals of Statistics*, vol. 38, no. 2, pp. 894–942, 2010.

[41] J. Liu, J. Jin, and Y. Gu, "Relation between exact and robust recovery for f-minimization: A topological viewpoint," in *Proc. IEEE Int. Symp. Inform. Theory*, 2013, pp. 859–863.

[42] Y. Zhang, "Theory of compressive sensing via ℓ_1 minimization: A non-RIP analysis and extensions," Technical report tr08-11 revised, Dept. of Computational and Applied Mathematics, Rice University.

[43] E. Candés and J. Romberg, "ℓ_1-magic: Recovery of sparse signals via convex programming," http://users.ece.gatech.edu/~justin/l1magic/, 2005.

[44] A. Belloni, V. Chernozhukov, and L. Wang, "Square-root lasso: pivotal recovery of sparse signals via conic programming," *Biometrika*, vol. 98, no. 4, pp. 791–806, 2011.

[45] Z. Ben-Haim, Y. Eldar, and M. Elad, "Coherence-based performance guarantees for estimating a sparse vector under random noise," *IEEE Trans. Signal Process.*, vol. 58, no. 10, pp. 5030–5043, 2010.

[46] R. Chartrand and W. Yin, "Iteratively reweighted algorithms for compressive sensing," in *IEEE Int. Conf. Acoust. Speech Signal Process.*, 2008, pp. 3869–3872.

[47] J. Fan and R. Li, "Variable selection via nonconcave penalized likelihood and its oracle properties," *Journal of the American statistical Association*, vol. 96, no. 456, pp. 1348–1360, 2001.

[48] A. Eftekhari, M. Babaie-Zadeh, C. Jutten, and H. Abrishami Moghaddam, "Robust-SL0 for stable sparse representation in noisy settings," in *IEEE Int. Conf. Acoust. Speech Signal Process.*, 2009, pp. 3433–3436.

[49] R. Chartrand, "Nonconvex splitting for regularized low-rank + sparse decomposition," *IEEE Trans. Signal Process.*, vol. 60, no. 11, pp. 5810–5819, 2012.

[50] R. Chartrand, "Fast algorithms for nonconvex compressive sensing: MRI reconstruction from very few data," in *IEEE International Symposium on Biomedical Imaging*, 2009, pp. 262–265.

[51] D. Malioutov and A. Aravkin, "Iterative log thresholding," in *IEEE Int. Conf. Acoust. Speech Signal Process.*, 2014.

[52] G. Gasso, A. Rakotomamonjy, and S. Canu, "Recovering sparse signals with a certain family of nonconvex penalties and dc programming," *IEEE Trans. Signal Process.*, vol. 57, no. 12, pp. 4686–4698, 2009.

[53] A. Miller, *Subset selection in regression*, CRC Press, 2002.

[54] W. Ortega, J. Rheinboldt, *Iterative solution of nonlinear equations in several variables*, vol. 30, SIAM, 2000.

[55] R. Rockafellar and R. Wets, *Variational analysis*, Springer, 1998.

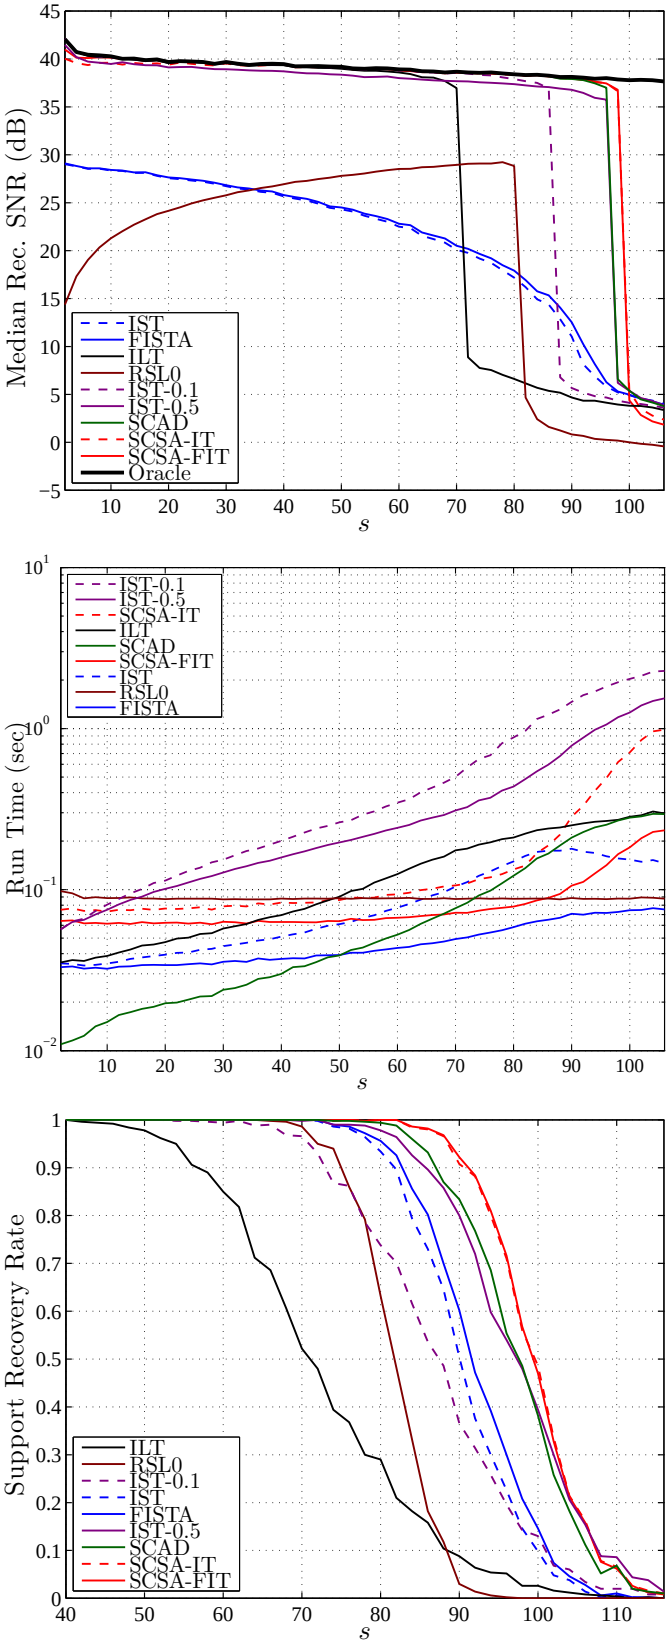


Fig. 6. Comparison of two variants of the SCSA algorithm in solving the noisy problem with settings similar to those of Fig. 5 except that the sparse vectors are Rademacher distributed. Trials are repeated 500 times, and results are averaged over them. Since SCSA-FIT and SCSA-IT as well as SCAD and IST-0.5 have very similar performances in terms of MSNR_{rec}, their traces are almost coincident in the top plot.